The Effect of Social Security Benefit Taxation on Beneficiary Income, Senior Labor Supply, and Program Financing

Jackson J. Howell*

May 1, 2024

Click here for updated manuscript.

Abstract

Social Security benefits first became taxable in 1984. Because benefit tax liability depends on level of current income, the tax may affect the current labor decisions of beneficiaries. I estimate its impact on labor supply and net Social Security finances. Unusually, the brackets are not indexed to inflation: I find that the proportion of benefits that are taxable has doubled and average tax rates on those benefits have tripled since 1994. Using earnings before benefits are claimed, I estimate counterfactual earnings in the absence of the benefit tax in order to identify treated beneficiaries. I find statistically significant, negative labor supply responses to the benefit tax among retired worker beneficiaries, who respond to a 1% change in their average tax rate with, on average, a 0.115% reduction in labor force participation and a 0.121% reduction in intensive-margin earnings. Using my estimates, I find that a $1 increase in benefit tax revenue is associated with a $0.26 reduction in payroll taxes. This research provides evidence on an increasingly important component of Social Security revenue and has implications for addressing the program's current financing shortfall.

*Yale University Tobin Center. howelljacksonj@gmail.com, https://sites.google.com/view/jacksonjhowell.

1 Introduction

In 1984, Social Security benefits first became subject to taxation, with up to 50% of benefits includable in taxable income; in 1994, up to 85% of benefits became taxable. The level of benefits included in taxable income depends on beneficiaries' current income, including earnings, meaning that the tax could affect those beneficiaries' current labor supply choices. The tax brackets which determine benefit tax liability were set in nominal terms and have *not* been indexed to inflation, so their real value has declined over time as prices have increased. While relatively few beneficiaries had taxable benefits in 1984, inflation has caused benefit taxes to become increasingly important for the finances both of beneficiaries themselves and of the program as a whole.

Social Security faces the depletion of its trust fund in 2033, at which time the program deficit will result in all beneficiaries facing a 23% cut in benefits (Board of Trustees (2023)) absent any policy change. Benefit taxes were first put in place when the Social Security system was facing a pressing financing crisis. When lawmakers reform the program to address the financing shortfall of today, the benefit tax might be revisited. Motivated by the need for program reform and evidence on this policy, this paper seeks to quantify the effects of the benefit tax in order to inform any proposed change to the policy.

This paper studies the effects of the benefit tax on beneficiary income, elderly labor supply, and Social Security finances as its brackets have crept lower and lower in real terms over time. After briefly providing institutional background and an explanation of how the taxes are computed, I measure and describe how benefit tax liabilities have increased over time across income groups. Then, I estimate labor supply responses to the benefit tax along the extensive and intensive margins. Finally, I estimate the effect of the benefit tax on net program financing, since a change in labor supply also changes Social Security's income from payroll taxation.

The first part the paper is institutional background. I explain how the benefit tax causes households in the middle of the income distribution to face extremely high marginal tax rates, rates that can even be higher than the maximum bracket rate. For instance, a single, 62-year-old beneficiary in 2023 with a total income of $65,000–of which $25,000 are Social Security benefits– would face a marginal tax rate of 40.7%, higher than the 37% statutory rate applied to incomes in excess of $578,125.[1] This fact motivates my investigation of behavioral responses to the tax, given that beneficiaries might respond to these extreme marginal tax rates.

Next, I measure the impact of the benefit tax over time in aggregate and broken down by income level. I find that the proportion of aggregate benefits includable in taxable income has risen from 4.5% in 1984 to 12.2% in 1994–when the second bracket was added–to 34.3% in 2019. Over the same period, I estimate that the average tax rate on benefits rose from 1.5% to 2.8% to 7.2%. The proportion of benefits subject to the tax has therefore tripled, and average tax rates doubled, without any legislative change to the policy since its 1994 revision. I document how these increases in tax exposure are driven by increasing liabilities among beneficiary households in the middle of the income distribution. I also show that benefit tax liability is rising among households in the bottom half of the income distribution. Therefore, this tax is becoming less progressive over time.

In the third part, I estimate the causal effect of the benefit tax on the labor supply choices of elderly beneficiaries. I use a methodology from Gelber et al. (2022), who study the labor supply effect of the Social Security earnings test. The earnings test reduces the benefits of working beneficiaries depending on how much earnings exceed a given threshold. Like the benefit tax, then, this policy could affect labor outcomes since after-tax benefit income changes depending on the beneficiary's level of earnings.[2] My strategy is to impute "potential earnings," or what an individual's earnings *would* have been in the absence of the benefit tax, using their real earnings in some base period before the benefit claiming decision. I use these potential earnings to identify the beneficiaries whose preferred level of income would subject them to the benefit tax.

I then use two-way fixed effects models to estimate the causal impact of being subject to the benefit tax on the probability of labor force participation and on earnings conditional on participation. This specification isolates the benefit taxation effect by controlling for changes in labor outcomes due to inherent differences between the treatment and control groups and due to the Social Security claiming decision. I estimate statistically significant, negative labor supply responses to the benefit tax along both margins. I estimate that an 1% increase in the average tax rate reduces labor supply by 0.115% and reduces earnings among workers by 0.121%. I show that my estimated elasticities, while high, are consistent with the literature on labor supply among

[1] These marginal tax rates are computed using NBER's TAXSIM model (see Feenberg and Coutts (1993)).

[2] Unlike the benefit tax, however, any benefit reduction from the earnings test will be compensated by increased benefits later in life.

marginally attached workers. See Chetty (2012) for a survey of evidence on intensive and extensive margin elasticities, including bounds on these elasticities with which my estimates are consistent.

Finally, I apply these elasticities to estimate how Social Security taxable payroll changes in response to the benefit tax. I estimate that every extra dollar of benefit tax revenue is accompanied by a $0.26 reduction in payroll tax revenue. This estimate characterizes the effect of the tax on net program finances, taking into consideration the large and negative labor supply elasticities with respect to the tax. I show through a back-of-the-envelope calculation that my elasticities of labor supply are consistent with this figure, and I show that this figure is consistent with the literature on deadweight loss to taxation.

Prior Literature There is a small literature that studies Social Security benefit taxation. In Goodman and Liebman (2008), the authors conceptualize the tax as a type of means testing, phasing out part of benefits as other-income rises. Burman et al. (2013) and Burman et al. (2014) investigate whether earnings bunch around the tax thresholds, which represent kinks in beneficiaries' budget constraints. In both papers, the authors do not find statistically significant evidence of bunching. Since these papers use bunching estimation, they are only able to estimate intensive-margin effects. I broaden this analysis to estimate effects along both margins using a TWFE framework.

One fact the previous papers point out that may be relevant for my research is that the tax may not be expected to have a strong effect on labor supply because it is not well understood by beneficiaries. In Greenwald et al. (2010), the authors find that only 57% of beneficiaries correctly answered that retirees' benefits are taxed when their income exceeds a certain threshold. Liebman and Luttmer (2015) connect how knowledge about Social Security rules can affect beneficiary labor supply by showing that beneficiaries change their labor supply after learning more about program provisions.

Page and Conway (2015) use the Current Population Survey to estimate the labor supply effects of the tax's introduction in 1984. They employ a difference-in-differences strategy to compare changes in labor supply among individuals treated by the tax (due to high income) against changes for individuals untreated by the tax (due to low income). They also exploit differences in state taxation of benefits. The authors focus on high income beneficiaries, for whom the tax reduces net income but does not affect the marginal tax rate of wages. Consistent with this group facing a pure income effect, the authors estimate that the tax increased labor force participation by two to five percent. A contribution of my paper is expanded analysis of beneficiaries at all parts of the income distribution, for whom the labor supply effect is ambiguous given competing income and substitution effects.

In Jones and Li (2018), the authors construct and fit a life-cycle model of labor supply, Social Security, and taxation. They find that replacing the benefit tax with a revenue-neutral increase in the payroll tax would increase aggregate labor supply, consumption, and welfare. The benefit tax has significant effects on the returns to work for the elderly, a group with high labor supply elasticities. The authors find that the benefit tax's labor supply effect operates primarily through the extensive margin. Interestingly, their model predicts that the welfare-distorting effects of the benefit tax is greater than that of the Social Security Earnings Test, which has received far more attention, both in the literature and politically.

This paper also fits into the literature studying the effects of Social Security policy on the work decisions of elderly beneficiaries. While Jones and Li (2018) predict large labor supply responses of beneficiaries to Social Security policies including the benefit tax, these have often not been empirically verifiable. Engelhardt and Kumar (2014) summarize the findings of the literature on the Social Security earnings test. These papers generally find statistically significant, but small, labor supply effects along the extensive and intensive margins. More recently, Gelber et al. (2021) find large and significant labor supply responses to the earnings test. My paper develops this literature by using a recent methodology to estimate effects on the entire income distribution, providing evidence that beneficiaries do respond strongly to Social Security policies as is theoretically predicted.

This paper proceeds as follows. Section 2 describes the context of the policy's introduction and how the benefit tax is calculated. Section 3 shows how the tax has evolved over time in aggregate and broken down by income groups. Section 4 describes the 2004 Benefits and Earnings Public-Use File, which I use to estimate labor supply effects of the benefit tax. Section 5 describes the empirical strategy I use to estimate those responses. Section 6 describes my estimates for the labor supply responses using multiple approaches; I also compute the effect of the benefit tax on net program finances. Section 7 concludes.

2 Institutional Background

Prior to 1984, Social Security benefits were not subject to federal or state income taxes. Recommendations by the 1979 Advisory Council and Greenspan Commission, convened during a time of financial stress to the Social Security system, recommended including up to 50% of benefits in taxable income.[3] Proceeds from these taxes are divided between the Old-Age and Survivors Insurance (OASI), Disability Insurance (DI), or National Railroad Retirement Investment (NRRI) trust funds. Beginning in 1994, up to 85% of benefits became taxable, with revenues from the next 35% of benefits earmarked for the Medicare (HI) trust fund. In Section 2.1, I describe how the benefit tax calculation can result in beneficiaries facing extremely high marginal tax rates. In Section 2.2, I provide a brief summary of the policy's history and the decision *not* to index the benefit tax brackets to inflation.

2.1 The benefit tax causes high marginal tax rates for beneficiaries in the middle of the income distribution

The amount of taxes an individual owes on their benefits depends on their "modified adjusted gross income" (MAGI), which is equal to adjusted gross income plus non-taxable interest income plus one-half of Social Security benefits. MAGI combines income and benefits from all members of the tax filing unit–including the head of household, spouse, and dependents. Benefits are "phased in" to taxable income based on how much MAGI exceeds certain thresholds, or brackets, which are different for single and joint filers.[4] Table A.1 presents the phase-in and phase-out thresholds for single and joint filers along with other key parameters used to calculate tax liability.

Figure A.1 shows how benefit taxes are computed for an example beneficiary, a single-filer with $6,000 in annual benefits in 1995.[5] Take their level of benefits as fixed, and consider how the beneficiary's tax liability changes as other income changes. In all panels, the X-axis represents the level of other income, and the Y-axis represents some key tax variable. For the rest of this discussion, I will refer to other income as "earnings," although it could in principle be non-labor, non-benefit income. This framework is useful not only to think through how the benefit tax works,

[3]See Dewitt (2001) for a more complete history of the policy's inception.

[4]See Internal Revenue Service (2022) for official rules governing the taxation of Social Security benefits.

[5]Such a benefit level is unlikely today, requiring someone to qualify for a "special minimum benefit," but would be more common in the past, and it allows us to see the four possible bend points in this process.

but also to contextualize how beneficiaries might choose their level of labor supply taking into consideration the tax implications on their *already determined* level of annual benefits.

Figure A.1a shows that benefits begin to be added to taxable income when earnings reach $22,000 (corresponding to $25,000 in MAGI). At this point, their phase-in rate is 50%, and every additional dollar of earnings adds $0.50 of benefits to taxable income. This lasts until the level of benefits in taxable income is 50% (see the left Y-axis), or $3,000 (see the right Y-axis). The second phase-in region begins when earnings reach $31,000 (and MAGI reaches $34,000). Then, every additional dollar of earnings adds $0.85 to taxable income. This lasts until the level of benefits in taxable income is 85% (see the left Y-axis), or $5,100 (see the right Y-axis).

Figure A.1b shows how taxable income (earnings plus taxable benefits) changes as earnings change. Outside of the phase-in regions, the slope of this line is 1: an additional dollar of earnings adds one dollar to taxable income. However, in the first phase-in region, the slope is 1.5: an additional dollar of earnings adds one dollar *plus 50 cents in benefits* to taxable income. In the second phase-in region, the slope is 1.85. The fact that the marginal increase in taxable income is greater than the marginal increase in gross earnings in these regions will be important for the next panel, showing the marginal tax rates faced by the beneficiary as their earnings increase.

In Figure A.1c, the solid orange line shows the effective marginal tax rate faced by the beneficiary at each level of earnings. Two things are immediately noticeable in this graph. First, marginal tax rates are non-monotonic. At the end of each phase-in region, marginal tax rates *decrease* as earnings increase. A decrease in marginal tax rates given an increase in income is, obviously, atypical in the United States's progressive tax system. Second, they face marginal tax rates that are unusually high given their place in the income distribution. In 1995, the highest federal income tax bracket rate was 39.6%, applied on income above $256,500 (about 513,183 in 2023 dollars) for single filers.[6] Meanwhile, if the beneficiary earned $27,500, resulting in $33,500 in total gross income, they would face a marginal tax rate of 42%. If they earned $32,500, they would face a marginal tax rate of 52%. That is, for individuals with an income equivalent to $55,020 today, their net income would increase by only 48 cents for every extra dollar earned.

As Burman et al. (2014) point out, the effective marginal tax rates faced by beneficiaries should

[6]In this paragraph, I use the BLS CPI Inflation Calculator to calculate real values for June 1995 in June 2023 dollars.

be thought of in relation to what they call the "underlying statutory rate," represented by the orange, dashed line in Figure A.1c. Every level of taxable income is associated with a marginal tax rate: typically, an extra dollar of income increases taxable income by one dollar, and the marginal tax liability is the marginal rate multiplied by $1. However, when benefits are involved, the relationship between earnings and taxable income is not one-to-one. In Figure A.1b, a one dollar increase in earnings can be associated with a $1.5 or $1.85 increase in taxable income. This extra taxable income (usually $1, but some multiple of $1 for the beneficiary) is multiplied by the corresponding marginal tax rate to get marginal tax liability. The underlying statutory rate is, in effect, multiplied by 1.5 or 1.85. This underlying statutory rate is the solid red line in Figure A.1c–during the phase-in regions, the beneficiary's effective rate is a multiple of the underlying rate. When phase-in ends, this multiplying effect also ends, and the effective marginal rates return to the underlying statutory rate, causing the non-monotonicity.

In summary, because benefits are phased into taxable income as earnings reach and exceed certain thresholds, gross taxable income increases by more than 1 dollar per extra dollar of earnings within these phase-in regions. Marginal tax rates are applied to this extra $1.50 or $1.85 in taxable income. The marginal tax liability is then the marginal rate multiplied by 1.5 or 1.85. This is why middle-income beneficiaries can face such high marginal tax rates. When phase-in ends, this multiplier effect stops, causing the drops in marginal tax rates as earnings increase.

2.2 The benefit tax brackets' real values have decreased over time

The thresholds which determine benefit tax liability (see Table A.1) were not indexed to inflation, meaning their real values decrease as prices increase. Over time, this means that more beneficiaries will have incomes exceeding the thresholds and will have taxable benefits. This *was* a conscious choice by policymakers, who intended more beneficiaries to become subject to the tax over time. However, this choice was not uncontroversial, as Senator William Armstrong of Colorado argued against it during legislative hearings:

> "...I do not like the notion that is abroad that somehow this is only going to be a tax on a handful of upper income, wealthy, affluent Social Security recipients, because in a very few years, unless our economic performance with respect to inflation is better than

I think it is going to be, everybody's benefits are going to be subject to taxation."[7]

Regardless, the brackets were not indexed to inflation, and when new brackets were added in the 1993 amendments, these also were fixed in nominal terms.

Figure A.2 shows how the real values of these benefit tax thresholds have decreased over time. Roughly, prices have tripled since 1984 and doubled since 1994–prices have increased by around 15% in the last two years alone. In 1984, when the first brackets of $25,000 ($32,000) for single (joint) filers were added, these were equivalent to $72,235 ($92,461) in 2023 dollars. When the second brackets of $34,000 ($44,000) were added, their equivalent value was $69,507 ($89,950). While the inflationary period of the 1970s that worried Senator Armstrong had ended by 1984, there has still been significant erosion to the real levels of income that would subject a beneficiary to taxation.

3 Benefit Taxation's Increasing Scope

3.1 Inflation has increased taxable benefits, the average tax rate on benefits, and the share of Social Security revenue from benefit taxes

In this section, I collect and combine data from multiple official sources and reports over the decades to calculate aggregate trends in benefit taxation. The aggregate level of taxable benefits, the aggregate tax liability owed on those taxable benefits, and the proportion of Social Security revenue from benefit taxes have increased over time. However, these data are not able to be broken down by level of beneficiary income.[8] I present statistics computed using these aggregated annual numbers here, and I defer discussion breaking down these liabilities by income level to Section 3.2.

Figure A.3 show aggregated, annual data that demonstrate the increasing importance of benefit taxation over time. Figure A.3a shows that the share of total benefits that are includable in taxable income has risen from 4.5% in 1984 to 12.2% in 1994 (when the second bracket was added) to 34.3% in 2019. The average tax rate on benefits also increased, as can be seen in Figure A.3b, from 1.5% in 1984 to 2.8% in 1994 to around 6% in 2014, the last year for which these numbers have been

[7] See Social Security Administration (1983) for the complete transcript of Senate hearings on the 1983 amendments.

[8] The IRS Statistics of Income Series (Internal Revenue Service (1984)) only provides breakdowns of taxable income by aggregated level of adjusted gross income. In addition, it does not show any beneficiaries who *did not* file a tax return in that year, making it hard to compute statistics on the entire population of beneficiaries. Furthermore, it does not provide any data on tax liability. The only data on benefit tax liability is published on an infrequent basis at the aggregated, annual level.

updated. Since 1994, there have been no changes to the benefit tax policy. All of these increases have been the result of "bracket creep," as the real value of the tax thresholds have decreased over time. This phenomenon has not been reflected in the larger tax system: from 1984 to 2019, the average total federal tax rate only slightly changed from 20.9% to 19.4% (Tax Policy Center (2023)).

The growing base for benefit taxation has made it an increasingly important source of revenue for the Social Security trust funds, as demonstrated in Figure A.3c.[9] At the policy's outset, less than 2% of income came from benefit taxes. By 2022, the benefit tax is nearly 4% of program income. These numbers are large in the context of the current financing shortfall: in 2022, the deficit would have been 45% larger if there were no revenues from benefit taxation. The present value of accumulated benefit tax revenue through 2022 is $1.2 trillion, equal to about 42% of the existing trust fund reserve balance. Therefore, the benefit tax has been an important policy for keeping the program solvent in recent years.

As Figure A.3c shows, eventually the benefit tax will account for upwards of 8% of program revenue, double what it is now, under the Trustees' intermediate assumptions. Without any income from benefit revenue going forward, the projected 75-year financing shortfall of $22.4 trillion would be about 28.9% bigger. Although benefit tax revenue is small relative to payroll tax revenue, it is a non-trivial part of Social Security financing over the coming decades. This makes it an important policy to consider during the next round of reforms to Social Security.

3.2 Benefit taxes have increased greatly for middle-income households

In order to calculate changes in liability broken down by income level, I use the CPS-ASEC survey. I use fields from the Census tax model to sort individuals into tax filing units and to calculate tax unit AGI and the level of benefits included in taxable income. Then, I use NBER's TAXSIM model (see Feenberg and Coutts (1993)) to estimate the liability owed on those benefits. See Appendix A for more details on my methodology, and there I also show evidence that my estimates are consistent with the aggregated data in figures presented in Section 3.1. My resulting dataset contains estimates on benefit tax liability (as well as total tax liability) for all survey households

[9] Only taxes levied on the first 50% of benefits go to the OASDI trust funds, with the rest going to the Medicare trust funds. This analysis does not include an analysis of the Medicare trust funds, however my findings are, in principle, applicable to that context.

that received Social Security benefits between 1991 and 2019.[10]

In order to investigate how the distribution of benefit tax liability has changed over time, I group individuals into their tax-filing unit, which I will refer to as households from now on. Each household is then grouped into an income bin based on their level of income relative to other households that year. I divide households into one group containing the bottom half of households, and five deciles splitting the top half of households.

Figure A.4 shows the evolution of benefit taxation among income groups. Figure A.4a shows the proportion of beneficiaries in each income group subject to the benefit tax–while virtually no beneficiaries with an income below the 70th percentile were taxed in 1991, by 2019 almost all beneficiaries in the top half of the income distribution are taxed. Indeed, even the bottom half of the distribution is becoming taxed at an accelerating rate in recent years.

Figure A.4b tells a similar story for the level of benefits included in taxable income: all income levels (except the top) have seen a rise in benefits subject to taxation. Households in the 8th income decile (between the 70th and 80th income percentiles) went from having almost no benefits taxed to having nearly 80% taxed in 2019. Figure A.4c shows the average tax rate on benefits by income bin. Again, beneficiary households between the 50th and 90th percentiles have seen an increase in their tax exposure over time due to the effects of inflation. Figure A.4d shows average tax rates on *non-benefit* income, demonstrating that this phenomenon is not reflected in the general tax code. While tax rates fluctuate with business cycles and tax reforms (particularly for the top decile), the continuous increase in average tax rates is not seen for non-benefit income.

These panels show that the increase in the benefit tax's reach have come from increasing exposure among middle and upper-middle income households. This pattern is unique to the benefit tax and is not reflective of some larger shift in the American tax system. It is also clear that this trend will continue, bringing in households below median income at an accelerating rate. As more Americans become subject to the benefit tax, it is becoming more important to understand the labor supply effects of the tax. This motivates my investigation into this question in the remainder of this paper.

[10] I begin analysis in 1991 because IPUMS only begins linking the tax model variables I need in the 1992 CPS-ASEC. Future work on this paper would involve manually linking prior, raw versions of the CPS-ASEC to my IPUMS extract to get the required fields.

4 Data

The empirical strategy I use to estimate labor supply effects of benefit taxation relies on observing individuals' earnings and benefit income over time. While the Current Population Survey tracks households for two years, observing units over longer periods will allow me to test for statistically significant pre-trends in labor supply variables between control and treatment groups. It will also allow me to test whether the treatment effect intensity changes over time. This is plausible, given the possibility that beneficiaries take time to detect and react to the tax.

As such, I use the 2004 Benefits and Earnings Public-Use File, or BEPUF (see Social Security Administration (2005)). This is the only public-use dataset made available by Social Security that contains individual-level, panel data on earnings and benefits in the period after 1994, when the benefit tax policy was last changed.[11] The dataset is a random sample of beneficiaries entitled to receive a benefit in December of 2004, linked to their monthly benefit amount, claim age, and longitudinal earnings histories from 1951 to 2003. Using the claiming information, monthly benefit, and historical cost-of-living adjustments, I calculate annual benefit levels each year since claim age. I limit the analysis to retired worker beneficiaries–excluding disabled workers as well as spouses, children, and survivors of retired worker beneficiaries.

The three key outcomes of interest, on which I estimate the effects of benefit taxation in Section 6, are labor force participation, earnings conditional on participation, and Social Security payroll taxes. Payroll taxes are easily computable in the BEPUF, since taxable earnings are provided as a field. However, earnings not in Social Security covered employment do not appear in the BEPUF. This complicates my estimation of labor force participation and intensive-margin earnings. I make a simple assumption that no beneficiaries in the BEPUF work in non-covered employment.

I also do not observe whether beneficiaries in the BEPUF are married or have dependents. As such, I cannot impute filing status (i.e., single or joint). This is non-ideal, since the benefit taxation thresholds depend on filing status. In the CPS dataset, a slight majority of beneficiary tax-filing units (around 58%) filed jointly since 1991. In the absence of a clearly better alternative, I simply assume that all beneficiaries file jointly and face the corresponding thresholds.

[11] The 2006 Earnings Public-Use File (see Compson (2011) for more details) is a similar dataset with some advantages over the BEPUF. However, it does not contain data on when benefits were claimed, making it impossible to impute exact annual benefits.

5 Empirical Strategy

5.1 Using prior earnings to impute counterfactual earnings in the absence of benefit taxation

Motivation In order to determine the effect of the benefit tax on labor supply, I must first identify which individuals are affected by it at all. It is not sufficient merely to observe which individuals owe taxes on their benefits. Some individuals may have participated in the labor force had it not been for the tax liability they would have incurred on their benefits. These individuals are obviously treated by the tax, but they will appear as untaxed in the data. Ideally, my data would contain the amount an individual would earn in the absence of the benefit tax. But, this counterfactual is not observed, and so it must be estimated. From here, I will refer to this level of counterfactual earnings as "potential" earnings.

I employ a methodology used by Gelber et al. (2022) to estimate potential earnings in the absence of benefit taxation. Their research investigates the labor supply effects of the Social Security earnings test, and I adapt their empirical strategy to the benefit taxation context. I assume that, in the absence of some shock, many individual's earnings in one year will be equal to their earnings in the next year after controlling for inflation.

Consider a group of 60-year-olds who will begin claiming benefits at age 62. Their potential earnings at age 62, taking age 60 as the base period, will be the real value of age 60 earnings at age 62. Potential MAGI (the key income variable used to determine whether benefits are taxable) is equal to their potential earnings plus one-half of annual benefits. Some of these individuals will have potential MAGI above the taxable threshold, and some of them will have potential MAGI below the threshold. I consider the first group to be "treated" by the benefit tax. This treatment imputation rests on the assumption that their real age 60 earnings are a good proxy for their age 62 earnings under no behavioral response to benefit taxation.

Of course, some individuals' preferred earnings will be different than my estimate of their potential earnings. This could be because of some shock to their productivity, wealth, or family structure. One could also reasonably expect that the benefit claiming decision would reduce work for many people. For individuals whose preferred earnings are lower than my estimate, I may falsely categorize them as treated. Conversely, I may incorrectly assign treated units to the control group if their preferred earnings are higher than my estimate. This causes measurement error in my treatment variable due to non-compliance of individuals "assigned" to treatment and control groups, but this would only make my estimates more conservative by biasing them towards zero. My measurement error will decrease as the share of individuals whose preferred real level of earnings is constant increases.

Formalization Let individuals be indexed by i and age by t. Each individual begins receiving Social Security benefits at claim age $c = c(i)$. The individual's annual benefits are $b_{i,t}$, and annual inflation is given by π_t.[12] Finally, $z_{i,t}$ represents Social Security taxable earnings. My goal is to estimate potential earnings in the absence of the benefit tax, denoted $\tilde{z}_{i,t}$. In order to do this, I use earnings in the base period, which I define to be two years before the claim age. For $t > c-2$, potential earnings are given by

$$\tilde{z}_{i,t} = z_{i,c-2} \cdot \Pi_{c-2}^{t-1} \pi_t, \tag{1}$$

which is simply the real value of base-period earnings. I take this to be the level of earnings an individual would choose if they exhibited no behavioral response to the benefit tax. The level of potential modified adjusted gross income (MAGI), denoted $\tilde{m}_{i,t}$, is

$$\tilde{m}_{i,t} = \tilde{z}_{i,t} + \frac{b_{i,t}}{2}. \tag{2}$$

The indicator for whether an individual's potential income subjects them to the benefit tax is then

$$\tilde{T}_{i,t} = \mathbb{1}(\tilde{m}_{i,t} \geq h(i)), \tag{3}$$

where $h(i)$ is the first benefit tax threshold, after which benefits begin being added to taxable income. This threshold varies by an individual's filing status. However, as discussed in Section 4, I assume all beneficiaries file jointly, and so this variable is always \$32,000. Note that I do not include any income in my potential MAGI formula other than earnings and Social Security benefits. This is necessitated by a lack of data on other income sources in the BEPUF. However, as with the potential income assumption in general, this assumption makes my treatment indicator noisy and

[12] These age-specific inflation factors are different across individuals, because the year in which one individual is a specific age may be different from the year in which another is that age.

will make my estimates more conservative.

In order to estimate tax effects on earnings conditional on participation, I define a variable equal to the ratio of actual and potential earnings. I take this variable to be the percent difference in amount worked for an individual relative to their base year, under the assumption that an individual's real wage is constant. This is an imperfect estimate, but as long as the measurement errors among the treatment and control groups evolve in parallel, the TWFE estimator will recover the causal effect on amount worked.

5.2 Estimating Treatment Effects using Two-Way Fixed Effects Models

I employ two-way fixed effects models to estimate the effects of benefit taxation on outcomes of interest–namely: labor force participation, intensive margin earnings, and Social Security payroll taxes. These specifications allow me to control for period-specific changes in the outcome variable shared between the treatment and control groups. It is likely that all of the outcome variables decrease significantly at the benefit claim decision. However, as long as the changes in these outcome variables would have been parallel in the absence of benefit taxation, the TWFE estimator will control for these shared changes and isolate the effect of the benefit tax. In Section 6.1, I will show evidence that trends are parallel in the pre-treatment period.

Static Two-Way Fixed Effects Model Treatment depends on several variables which differ by individual. Since treatment adoption is staggered, I abstract from the two-period difference-in-differences model and employ a static two-way fixed effects (TWFE) strategy as my most basic estimator to identify the effect of the benefit tax on labor outcomes. The static TWFE model is given by

$$Y_{i,t} = \alpha_i + \phi_p + \beta \tilde{T}_{i,t} + \gamma_t + \delta_y + \text{sex}_i + \log(b_{i,t}) + \epsilon_{i,t}. \tag{4}$$

$Y_{i,t}$ denotes outcomes of interest, and the TWFE estimator is β. This specification controls for time-invariant differences in outcomes between individuals as well as period-specific shocks. Period (denoted p) gives the time relative to the benefit claiming decision and equals 0 at claim age. I also use age t, year y, and sex fixed effects as additional controls, along with the natural logarithm of Social Security benefits.

The TWFE estimator is interpretable when treatment effects are homogeneous across time (see Roth et al. (2023) for a survey of the recent literature on TWFE and generalizations that address its shortcomings). This assumption is unlikely to hold in my setting, both because my noisy potential treatment variable might pick up actually treated individuals as time goes on and because individuals may take time to detect and react to the benefit tax.

Dynamic Two-Way Fixed Effects Model The dynamic TWFE model addresses this problem. I define the variable $R_{i,t} = t - t'$, where t' equals the time of treatment. This variable gives the time relative to treatment and equals 0 in the first period of treatment. For individuals in the control group, $R_{i,t} = -1$. The dynamic TWFE model is

$$Y_{i,t} = \alpha_i + \phi_p + \sum_{t' \neq -1} \mathbb{1}(R_{i,t'})\beta_{t'} + \gamma_t + \delta_y + \text{sex}_i + \log(b_{i,t}) + \epsilon_{i,t}. \tag{5}$$

This specification allows for treatment heterogeneity with respect to time relative to treatment. The omitted period is the age before treatment.

Assumptions The TWFE models require typical difference-in-differences assumptions. Critically, trends in the outcome variable should be parallel between the treatment and control groups in the pre-period *and* should have been parallel in the post period in the absence of treatment. In order to make the treatment and control groups more comparable, I restrict my sample to individuals with potential MAGI close to the benefit tax threshold. First, I only keep observations from three years before treatment through two years after treatment. My income restriction is that potential MAGI in the first period must be at least \$22,000 (the threshold minus \$10,000) and potential MAGI in the last period must be less than \$42,000 (the threshold plus \$10,000).

Using the dynamic specification, I provide evidence on the parallel trends assumption in Section 6.1. I have already mentioned that the static TWFE estimator relies on the assumption that treatment effects are homogeneous with respect to time-to-treatment, but I relax this assumption in the dynamic TWFE model. However, the dynamic TWFE itself relies on an assumption that period-specific treatment effects are homogeneous across treated individuals. Future work on this paper would involve relaxing this assumption by introducing more novel event study estimators, such as the one proposed by Callaway and Sant'Anna (2021), that allow for treatment heterogeneity across units. An additional advantage of this specification is the ability to test for how treatment effects vary by age, sex, and income level.

6 Results

6.1 Benefit taxation reduces labor supply on extensive and intensive margins

I now present results on the relationship between benefit taxation and labor outcomes using the models specified in Section 5. In Table A.2, I show estimated coefficients for the effect of benefit taxation on labor force participation and earnings conditional on participation. Coefficients relating to payroll taxes are addressed in Section 6.3. Treatment here is defined when an individual has any amount of income above the tax threshold, and so these results reflect the average effects of *any* level of benefit taxation. Beneath each outcome variable, the first columns in Table A.2 show results from the static TWFE model. Estimates for the tax's effect on both margins are negative and statistically significant: the coefficients imply that the average effects of having potential income above the taxable threshold are a 2.2% reduction in the probability of participating in the labor force and a 2.5% reduction in earnings conditional on participation.

I also present results using the dynamic TWFE model, a specification which allows for the intensity of treatment to change as an individual is taxed for a longer period of time. Results are shown in the second column beneath each outcome variable in Table A.2, and these coefficients (along with 95% confidence intervals) are graphed in Figure A.5. Estimates on the intensive margin show evidence that negative treatment effects increase in magnitude over time, ranging from -1.9% at time of treatment to -5.2% two years after treatment. However, there is no evidence of changing intensity for the extensive margin, with the treatment effect hovering around -1.8% for the entire post-period. For both outcomes, there is no statistical evidence of pre-trends.

Heterogeneous effects with respect to time-from-treatment may be a result of beneficiaries becoming aware of the benefit tax over time. This is reasonable given little knowledge of the tax, at least until a beneficiary starts having to report it on a tax return. It is also possible that these effects are a result of treatment group compliance increasing over time. Given the possibility that I overestimate some individuals' potential earnings in the first period of treatment, some "treated" individuals may not actually be so. However, the year after, perhaps some of these individuals *actually* become treated. If the increase in compliance among the treatment group exceeds an increase in noncompliance in the control group, this would reduce the noise in my treatment proxy biasing results towards zero. I choose to use my most conservative estimates of the treatment effect in the first year of treatment as my preferred specification.

However, these estimates of the effect of *any* taxation are not easily interpretable, since some individuals who are taxed may face low average tax rates (as seen in Figure A.4). To address this problem, I convert these TWFE estimators into elasticities with respect to an increase in average tax rate in the next subsection.

6.2 Elasticities are large and consistent with marginal attachment to labor force

I would like to express the labor supply effects of benefit taxation as a percentage of the average tax rate: i.e., if the average tax rate increases by 1% due to benefit taxation, what is the average effect on their labor force participation or intensive-margin earnings? I define a new variable, the average tax rate, which is the fraction of *observed* income that would be taxed if an individual earned their *potential* income:

$$ATR_{i,t} = \frac{\tau(\tilde{z}_{i,t} + b_{i,t})}{z_{i,t} + b_{i,t}} \tag{6}$$

Here, $\tau(\cdot)$ is the tax function which takes income as an input and outputs tax liability. I identify the effect of benefit taxes on this average tax rate using the static and dynamic TWFE models from above. The TWFE estimator isolates the effect of the benefit tax on total tax liability expressed as a percentage of observed gross income. The estimator's numerator includes benefit taxes accrued on actual income and the benefit taxes accrued on the part of potential earnings which is not actually earned, because I want to estimate what the benefit tax liability would be before any behavioral response. It is also helpful here to define another variable, the average net-of-tax rate, which is just 1 minus the average tax rate: $ANTR_{i,t} = 1 - ATR_{i,t}$. I provide regression results for these two outcome variables using the static and dynamic models in Table A.4 and Figure A.8. The regression results are significant and in the expected directions, and there is no statistically significant pre-trend in the dynamic setting for either measure.

Dividing the TWFE estimator for labor force participation by the sample mean of labor force participation in the control group is the percent change in participation given the benefit tax. I can also calculate the percent change in average tax rate due to the benefit tax in the same way. Dividing these two percent changes gives me the elasticity of labor force participation with respect to 1% change in average tax rate on gross income:

$$\eta^E = \frac{\beta^{LFP}/\mu^{LFP}}{\beta^{ATR}/\mu^{ATR}}. \tag{7}$$

In this equation, the β terms represent regression coefficients of the given variable, and the μ terms represent the sample mean of the variable in the control group. η^E is the extensive margin elasticity, but the analogous intensive margin elasticity η^I can be calculated using results from the earnings-conditional-on-employment TWFE estimate. This process is also used to calculate elasticities with respect to the average net-of-tax rate.

Elasticities are presented in Table A.3. Using the static TWFE model, my estimated extensive and intensive margin elasticities with respect to the average tax rate are -0.069 and -0.087, respectively. More explicitly, these estimates suggest that the response to a 1% increase in the average tax rate is, on average, a 0.069% reduction in labor force participation and an 0.087% decrease in earnings conditional on participation.

The extensive (intensive) elasticity with respect to the average net-of-tax rate is 0.33 (0.413) in the static model. These estimates are consistent with the literature. The extensive elasticity is similar to that found by Gelber et al. (2022) in response to the Social Security earnings test and, as the authors point out, falls in the middle of the extensive margin elasticities in Chetty (2012), which are bounded by 0.15 and 0.43. My estimate for the intensive margin elasticity is also within 0.28 and 0.54, the bounds on (Hicksian) intensive margin elasticities suggested by Chetty. Among the studies surveyed, my estimates are most consistent with papers studying groups with potentially high labor supply elasticities–such as women or extreme high earners. This also makes sense for my context, as beneficiaries have a non-labor source of income they can resort to if they choose to reduce labor earnings.

My estimates are not easily comparable to the benefit tax effects found by Page and Conway (2015), since they study the labor supply effects of high-income individuals for whom all benefits are taxed. For this group, the tax has a pure income effect, and the author's estimated elasticities are positive. My treatment group, however, contains individuals who are near the threshold at which benefits begin being taxed. Therefore, my treatment effect being negative implies that the substitution effect outweighs the income effect, which is consistent with beneficiaries responding to the high marginal tax rates caused by the benefit tax near the threshold.

I also compute elasticities implied by the dynamic TWFE model, shown in Table A.3. Intensive margin elasticities are relatively stable across time to treatment, ranging from -0.099 to -0.121 with respect to the ATR and from 0.47 to 0.576 with respect to the ANTR. Extensive margin elasticities, however, decrease in magnitude over time across all the estimates, from -0.115 to -0.034 with respect to the ATR and from 0.548 to 0.163 with respect to the ANTR. As with the regression results, I use the estimated dynamic effect in the first year of treatment as my preferred specification.

The extensive (intensive) elasticity with respect to the average net-of-tax rate is 0.330 (0.413) in the static model. The extensive elasticity is similar to that found by Gelber et al. (2022) in response to the Social Security earnings test and falls in the middle of the extensive margin elasticities in Chetty (2012), which are bounded by 0.15 and 0.43. My estimate for the intensive margin elasticity is also within 0.28 and 0.54, the bounds on (Hicksian) intensive margin elasticities suggested by Chetty. Among the studies surveyed, my estimates are most consistent with papers studying groups with potentially high labor supply elasticities–such as women or extreme high earners. This makes sense for my context, as beneficiaries have a non-labor source of income they can resort to if they choose to reduce labor earnings.

6.3 Benefit taxes reduce Social Security's income from payroll taxes

I now present results using the static and dynamic TWFE models on a new outcome: Social Security payroll taxes. However, here I calculate elasticities of Social Security payroll revenue given an increase in benefit tax revenue. That is, if benefit tax revenue increases by $1, how much will payroll tax revenue decrease? This will indicate the net effect of the benefit tax on program finances given the previous evidence that beneficiaries reduce labor supply in response to the benefit tax.

Regression results using the static TWFE model on Social Security payroll taxes are presented in Table A.2. Results are negative and statistically significant, with the effect of *any* taxation being, on average, a reduction of $455. Results from the dynamic model are also in the table and graphed in Figure A.5c. The coefficient grows in magnitude, becoming more negative as time-relative-to-treatment increases. Although there is not a clear pre-trend, the estimated coefficient two periods before treatment is statistically significant. This problem merits further investigation.

As before, these coefficients are not easily interpretable because treated individuals may face very different amounts of benefit taxation. I convert these estimates to elasticities in a slightly different way than before. Here, I want to estimate the reduction in payroll taxes given a $1 increase in benefit taxation. The numerator of this elasticity is the TWFE estimator of the effect of benefit taxation on payroll taxes. The denominator is the mean amount of benefit taxes faced by treated beneficiaries.

Estimated elasticities of payroll taxation with respect to the benefit tax are presented in Table A.3. In the static model, this elasticity is rather high, at -0.59, implying a $1 increase in benefit tax revenue decreases payroll tax revenue by -$0.59. In the dynamic model, the elasticity at time of treatment starts at -$0.26 and decreases rapidly. As before, I take this value from the dynamic model at time of treatment as my preferred estimate for the effect of benefit taxation on payroll tax revenues.

I conduct a back-of-the-envelope exercise showing the labor supply elasticities in Section 6.2 are consistent with this estimate of deadweight loss. Among individuals who will be treated by the benefit tax, average payroll taxes are $3,719.25, average gross income is $26,140.2, and the average tax rate is 17.32%. A one percent change in the average tax rate is 0.173% of gross income, or $45.26 per person on average. From the previous section, the extensive and intensive margin responses to this tax shock are -0.115 and -0.121. The new average payroll tax can be calculated as a weighted average of the new average payroll taxes owed by the group of individuals who keep working and the group that leaves the labor force.[13] The group who leaves the labor force, weighted by 0.00115, have average taxes of $0, and the group who keeps working, weighted by $1 - 0.00115$, have average taxes of $(1 - 0.00121) \cdot \$3719.25$. The new average payroll tax comes out to be $3,710.47, meaning that average payroll taxes decreased by $8.77 on average. Therefore, for every dollar of benefit taxes gained, there were $0.19 in payroll taxes lost. While this estimate is not exactly equal to the above elasticity, it shows that the magnitude is reasonable.

[13]The payroll tax is a flat tax on taxable earnings, so I do not need to worry about the marginal tax rate on earnings decreasing as earnings decrease.

Estimates of the deadweight loss of taxation can range widely in the literature. Saez et al. (2012) suggests a baseline estimate of 27.7% among high-income earners. Giertz (2009) shows how the estimate of deadweight loss is highly sensitive to the elasticity of taxable income. When using an elasticity in the middle of estimates found in the literature, the figure is between 18% and 42% for all workers. Feldstein (1999) suggests the deadweight loss of income taxation to be as high as 32% on average. My estimate therefore fits in well with the magnitude of effects found in the larger literature.

7 Conclusion

The Social Security benefit tax, soon to turn 40-years-old, has not received much attention in the economics literature. This is partially explained by its relative lack of importance in its early years. A novel contribution of this paper is my demonstration that the benefit tax has evolved from a policy affecting only the highest-income beneficiaries to one that now affects over half of households receiving Social Security benefits. The benefit tax reduces net benefit income (and thus, total net income) as earnings rise, making it possible the tax affects labor supply. Unlike previous papers like Burman et al. (2013) and Burman et al. (2014), I use two-way fixed effects models to estimate labor supply effects instead of bunching estimation. This allows me to estimate extensive-margin responses as well as intensive-margin ones. While Page and Conway (2015) use a difference-in-differences framework, they focus on high-income individuals. I broaden my analysis to all taxed beneficiaries, since the labor supply responses of the households in the middle of the income distribution become more interesting as that group becomes more taxed over time. In order to do this, I adapt a methodology from Gelber et al. (2022) to identify the treatment group by estimating earnings in the absence of a behavioral response to the benefit tax. I conclude by asking a question, to my knowledge novel: given a labor supply response to benefit taxation, how much does an additional dollar of benefit tax revenue change Social Security's revenue from payroll taxation? My finding of a negative labor supply response implies that Social Security makes a trade-off between payroll taxes and benefit taxes.

Future work on this topic could include applying this methodology to restricted-use, administrative data, which would allow for new questions and could address the limitations of this paper.

For instance, a sample of Social Security's Master Earnings File linked to the Master Beneficiary Record (the data build used by Gelber et al. (2020)) would allow for analysis of the population of all beneficiaries since 1984, instead of just those receiving benefits in 2004. The greater sample sizes in administrative data would also allow for more narrow restrictions on income, making the control and treatment groups more comparable. Since the threshold around which these treatment and control groups are defined has decreased in real terms over time, one could even estimate elasticity heterogeneity by income. If the above data product could also link to survey data such as the CPS, the imputation of filing status and total household income (and therefore of treatment and control groups) would be more exact. My paper, which provides evidence of large labor supply responses, motivates future investigation into this area in the near future to provide further evidence on this little-studied policy as the need for Social Security reform intensifies.

References

Board of Trustees (2023). The 2023 OASDI Trustees Report. Technical report, Social Security Administration.

Burman, L., T. P. Center, N. B. Coe, and L. Tian (2013). Older Taxpayers' Response to Taxation of Social Security Benefits. In *Tax Administration at the Centennial: An IRS-TPC Research Conference*, pp. 3.

Burman, L. E., N. B. Coe, K. Pierce, and L. Tian (2014, June). The effects of the taxation of social security benefits on older workers' income and claiming decisions. *National Tax Journal 67*(2), 459–485. Publisher: The University of Chicago Press.

Callaway, B. and P. H. C. Sant'Anna (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Chetty, R. (2012). Bounds on Elasticities with Optimization Frictions: a Synthesis of Micro and Macro Evidence. *Econometrica 80*(3), 969–1018. Publisher: [Wiley, Econometric Society].

Compson, M. (2011). The 2006 Earnings Public-Use Microdata File: an Introduction. Technical Report 71(4).

Davies, P. S. and T. L. Fisher (2009). Measurement Issues Associated with Using Survey Data Matched with Administrative Data from the Social Security Administration. *Social Security Bulletin 69*(2), 1–12.

Dewitt, L. (2001, February). Taxation of Social Security Benefits. Technical Report 12, Social Security Administration.

Engelhardt, G. V. and A. Kumar (2014, June). Taxes and the Labor Supply of Older Americans: Recent Evidence from the Social Security Earnings Test. *National Tax Journal 67*(2), 443–458.

Feenberg, D. and E. Coutts (1993). An introduction to the TAXSIM model. *Journal of Policy Analysis and management 12*(1), 189–194. Publisher: JSTOR.

Feldstein, M. (1999). Tax Avoidance and the Deadweight Loss of the Income Tax. *The Review of Economics and Statistics 81*(4), 674–680. Publisher: The MIT Press.

Flood, S., M. King, R. Rodgers, S. Ruggles, J. R. Warren, B. Daniel, A. Chen, G. Cooper, S. Richards, M. Schouweiler, and M. Westberry (2023). CPS-ASEC 1985-2022.

Gelber, A., D. Jones, D. W. Sacks, and J. Song (2022). The Employment Effects of the Social Security Earnings Test. *Journal of Human Resources 57*(2), 341–371. Publisher: University of Wisconsin Press.

Gelber, A. M., D. Jones, and D. W. Sacks (2020). Estimating Adjustment Frictions Using Nonlinear Budget Sets: Method and Evidence from the Earnings Test. *American Economic Journal: Applied Economics 12*(1), pp. 1–31. Publisher: American Economic Association.

Gelber, A. M., D. Jones, D. W. Sacks, and J. Song (2021, October). Using Nonlinear Budget Sets to Estimate Extensive Margin Responses: Method and Evidence from the Earnings Test. *American Economic Journal: Applied Economics 13*(4), 150–93.

Giertz, S. H. (2009). The Elasticity of Taxable Income: Influences on Economic Efficiency and Tax Revenues, and Implications for Tax Policy. *Tax Policy Lessons from the 2000s*, 101.

Goodman, S. and J. Liebman (2008). The Taxation of Social Security Benefits as an Approach to Means Testing. In *Annual Conference of the Retirement Research Consortium*.

Greenwald, M., A. Kapteyn, O. Mitchell, and L. Schneider (2010, October). What Do People Know About Social Security. Technical Report WR-792-SSA, RAND Corporation.

Iams, H. M. and P. J. Purcell (2013). Social Security income measurement in two surveys. *Social Security Bulletin 73*(3), 1–10.

Internal Revenue Service (1984). Individual Complete Report. Technical Report Publication 1304.

Internal Revenue Service (2022). Social Security and Equivalent Railroad Retirement Benefits. Technical Report Publication 915, Internal Revenue Service.

Jones, J. B. and Y. Li (2018, March). The effects of collecting income taxes on Social Security benefits. *Journal of Public Economics 159*, 128–145.

Liebman, J. B. and E. F. P. Luttmer (2015, February). Would People Behave Differently If They Better Understood Social Security? Evidence from a Field Experiment. *American Economic Journal: Economic Policy 7*(1), 275–299.

Lin, D. (2022). Methods and assumptions of the cps asec tax model. *US Census Bureau*.

Martin, P. P. (2016, January). Why Researchers Now Rely on Surveys for Race Data on OASDI and SSI Programs: A Comparison of Four Major Surveys. *Research and Statistics Note 2016*(1).

Meyer, B. D., D. Wu, G. Finley, P. Langetieg, C. Medalia, M. Payne, and A. Plumley (2020, December). The Accuracy of Tax Imputations: Estimating Tax Liabilities and Credits Using Linked Survey and Administrative Data. In *Measuring Distribution and Mobility of Income and Wealth*, pp. 459–498. University of Chicago Press.

Page, T. F. and K. S. Conway (2015, May). The Labor Supply Effects of Taxing Social Security Benefits. *Public Finance Review 43*(3), 291–323. Publisher: SAGE Publications Inc.

Roth, J., P. H. C. Sant'Anna, A. Bilinski, and J. Poe (2023, August). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics 235*(2), 2218–2244.

Saez, E., J. Slemrod, and S. H. Giertz (2012, March). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature 50*(1), 3–50.

Social Security Administration (1983, March). Social Security Act Amendments of 1983. Report of the Committee on Finance, to Accompany S. 1. Technical Report 98-23.

Social Security Administration (1984). Annual Statistical Supplement. Technical report.

Social Security Administration (2005, October). Benefits and Earnings Public-Use File, 2004.

Tax Policy Center (2023). Historical Average Federal Tax Rates for All Households (1979 - 2019).

US Department of the Treasury (1984). Report on the Taxation of Social Security and Railroad Retirement Benefits. Technical report.

A Estimating Benefit Taxes using the Current Population Survey

I use extracts of the CPS Annual Social and Economic Supplement (CPS-ASEC) over the period 1984-2021 from IPUMS (Flood et al. (2023)). The CPS-ASEC contains data on demographic information and many income sources, including labor earnings, earnings from government tax and transfer programs, income from dividends and interest, and income from Social Security or private pensions. It is well known that Social Security data are subject to measurement error in these surveys, since they rely on self-reported information. See Davies and Fisher (2009), Iams and Purcell (2013), and Martin (2016) for descriptions of this problem. Figure A.6 shows how CPS estimates of aggregate benefits, number of beneficiaries, and average annual benefits differ from official SSA values by year over my period of 1984-2021.

The CPS persistently underestimates aggregate annual benefits and the number of beneficiaries by year, with an average percent difference from SSA sources of -15.2% and -14.5%, respectively. By contrast, the CPS's estimate of average annual benefits is much closer, with an average percent difference of -2.5%. In Iams and Purcell (2013), the authors use respondents' Social Security records to match them to their actual, administrative data. The authors find that, on average, the CPS underestimates the number of beneficiaries over the age of 60 by 9.6% and average annual benefits by 1.2% in 2008. The CPS extract, which includes beneficiaries of all ages, has corresponding underestimates of number of beneficiaries and average benefits of 16.2% and 1.2%, respectively, in 2008. This suggests that I am undercounting retired beneficiaries to some degree, but I am especially undercounting beneficiaries below the age of 60–either disabled, secondary, or surviving beneficiaries. However, average benefits are consistent with the idea that respondents who do report being a beneficiary are reporting fairly accurate benefit amounts.

I adapt a methodology from Meyer et al. (2020) to impute tax-filing units in the CPS. To identify filing status (i.e., primary filers, married individuals filing jointly, non-filers, etc.), I use variables from the CPS-ASEC tax model (see Lin (2022) for details on how this model's imputation of filing units has changed over time). Since these tax variables are only available in the IPUMS CPS datasets beginning in 1992, I limit this part of the analysis to the period 1992-2021. I use CPS variables linking children to parents to link any under-18 individuals in a household to their parental filing unit. The Census tax model imputes adjusted gross income for each filing unit, which may include Social Security income. Given this AGI and the level of tax-unit benefits, I solve for the level of benefits in taxable income using numerical methods. Figures A.7a and A.7b show how my estimates compare to official, aggregate data in the IRS's Statistics of Income series. My estimates are reasonably close given the use of survey data, with a mean percent difference of close to 0%.

For each tax unit, I convert CPS income fields into variables that can be input into NBER's TAXSIM model (see Feenberg and Coutts (1993)). I construct earnings, dividend income, interest income, income from government transfers, Social Security benefits, and "other property income." Other property income is constrained to be the difference between CPS-imputed AGI and all these other income sources. I calculate the total tax liability owed by each tax unit in actuality and the tax liability if their Social Security benefit income were zero. The difference between these figures is defined as the tax liability owed on Social Security benefits.

The bottom two panels of Figure A.7 show my estimates of average effective tax rates against official SSA numbers. My estimates are similar to the aggregate numbers, with a mean percent difference of just over 4%. Therefore, my imputation of benefit taxes in the CPS is largely consistent with official aggregate data.

Figure A.1: Benefit Taxation Example

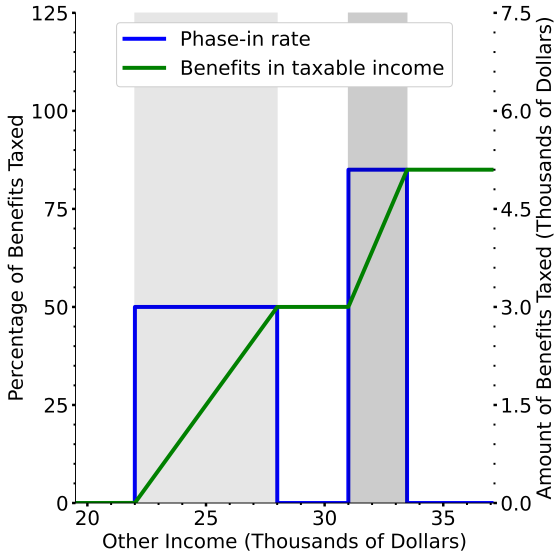

(a) Benefits in Taxable Income

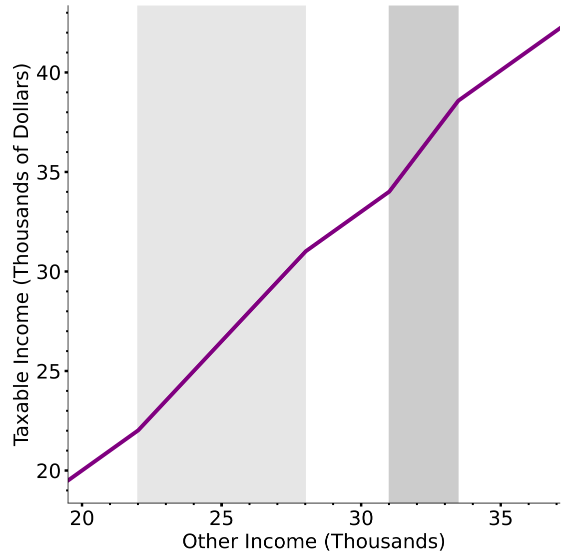

(b) Taxable Income

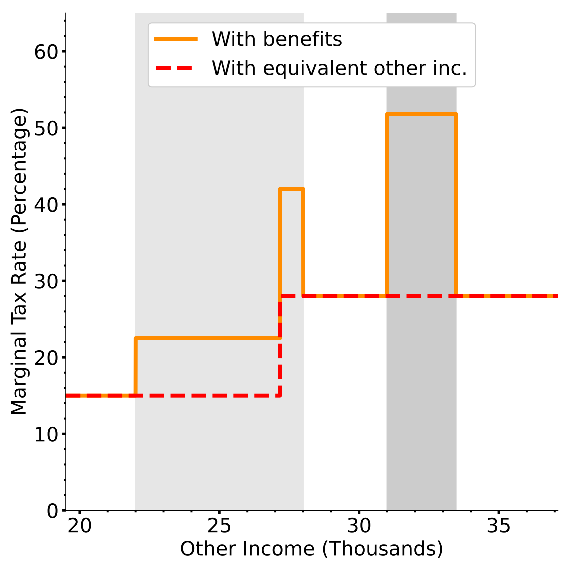

(c) Marginal Income Tax Rate

Notes: This figure shows how benefit taxes are calculated for a single-filer with benefit income of $6,000 as their level of other income changes. Figure A.1a shows how benefits are "phased in" to taxable income as other income exceeds certain thresholds. In the first (second) phase-in region, every extra dollar of other income adds $0.50 ($0.85) to taxable income. Figure A.1b shows this level of taxable income, equal to other income plus taxable benefits. Finally, in Figure A.1c, the orange line shows the marginal tax rates faced by the beneficiary as other income changes. The red dashed line shows the marginal taxes faced by an individual with the same level of taxable income, but no benefit income. In the phase-in regions, every extra dollar of income the beneficiary earns adds either $1.50 or $1.85 to taxable income (see the slope of the line in Figure A.1b). This extra amount is then multiplied by the underlying marginal rate to get marginal tax liability. Therefore, the beneficiary's effective marginal tax rates are equal to the underlying marginal rate multiplied by 1.5 in the first phase-in region and by 1.85 in the second. This phenomenon accounts for the unusually high marginal rates faced by the middle-income beneficiary and the non-monotonicity of their marginal rates at the end of the phase-in regions.

Figure A.2: Real Social Security Benefit Taxation Thresholds

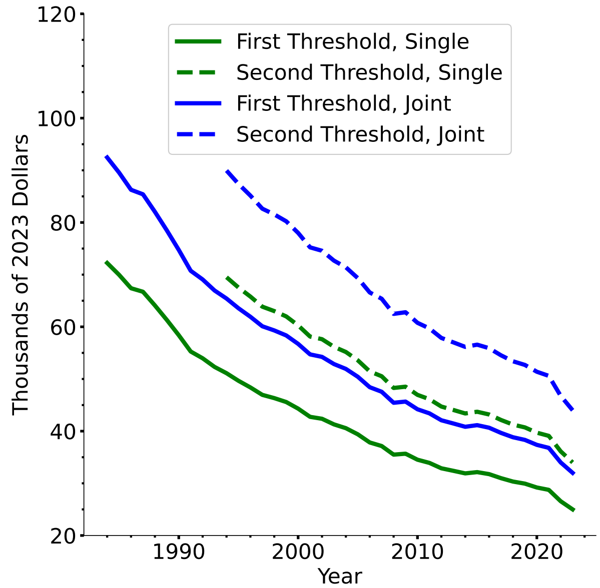

Notes: Y-axis units are real 2023 dollars, calculated using the January-to-January changes in CPI-W. Thresholds are the levels of modified adjusted gross income above which benefits are added to taxable income. These thresholds vary by filing status. The first set of thresholds were added in 1984, while the second set were added in 1994. See Section 2.1 for more details.

Figure A.3: Aggregate Trends in Benefit Taxation

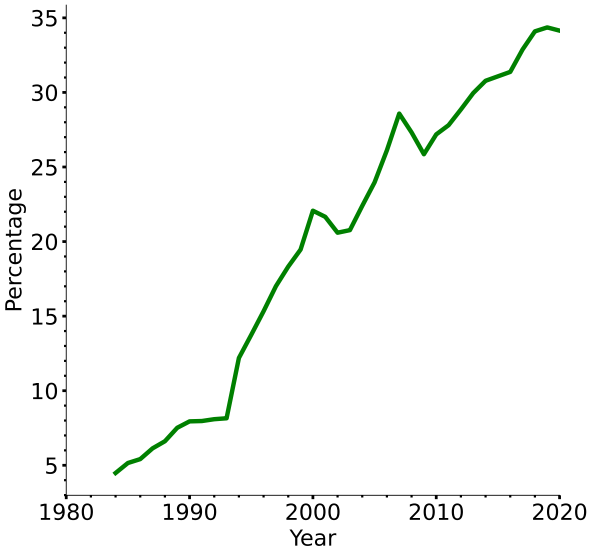

(a) Percent of Benefits Taxable

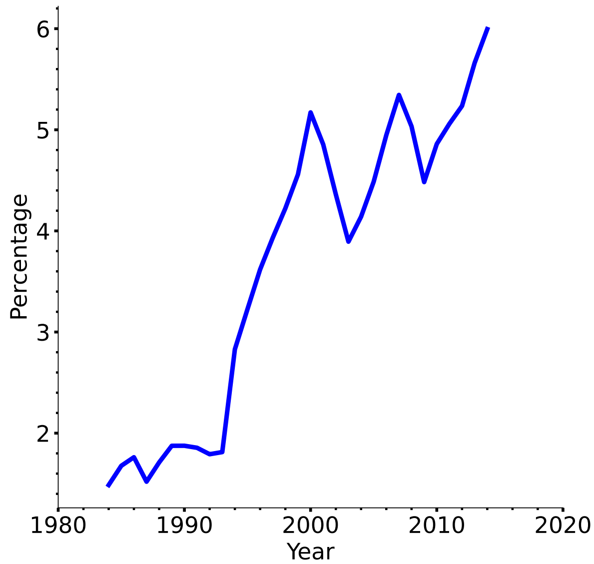

(b) Average Tax Rate on Benefits

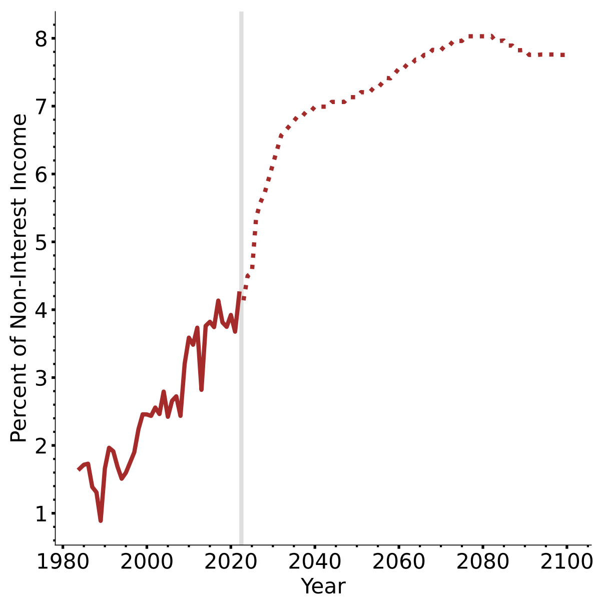

(c) Share of OASDI Revenue from Benefit Taxes

Notes: This figure shows how benefit taxation has become more important over time. Data are presented at the aggregated, annual level. Figure A.3a shows the average percentage, across all beneficiaries, of benefits that are includable in taxable income (i.e., the aggregate dollar amount of benefits includable in taxable income divided by the aggregate dollar amount of total benefits), using data from Internal Revenue Service (2022) and Social Security Administration (1984). Figure A.3b shows the average tax rate on benefits, equal to the total benefit tax liability divided by the total level of benefits, incorporating data from US Department of the Treasury (1984). This series has not been updated to include data after 2014. Figure A.3c shows the share of total Social Security income that is from benefit taxation, including projections from the intermediate, current-law assumptions of the 2023 Trustees Report (Board of Trustees (2023)). Projected data are represented by the dotted line and come after 2022, which is marked by the vertical gray line.

Figure A.4: Trends in Benefit Taxation by Income Group

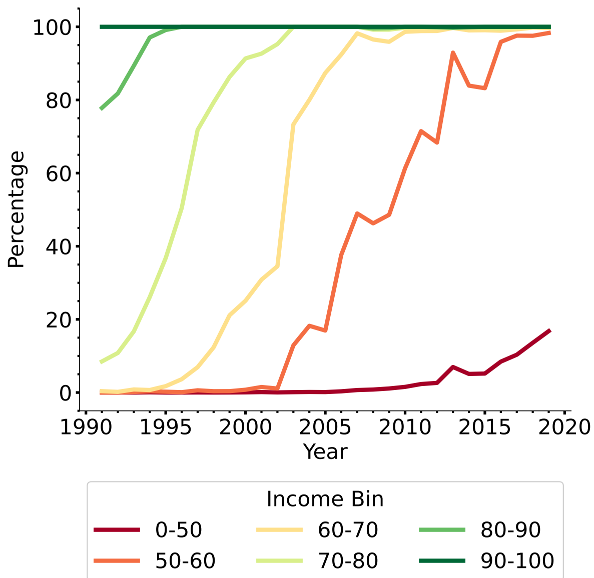

(a) Percent of Beneficiaries Taxed

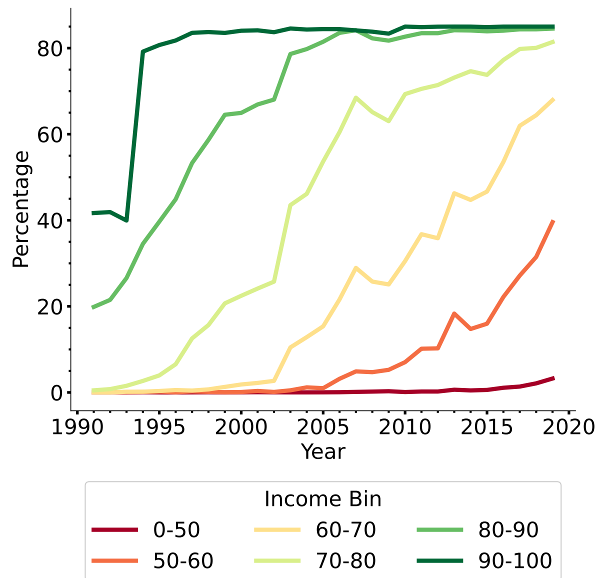

(b) Percent of Benefits in Taxable Income

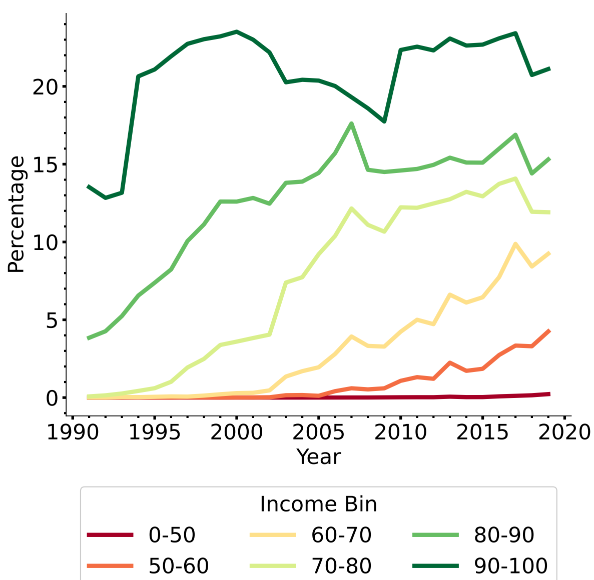

(c) Average Tax Rate on Benefits

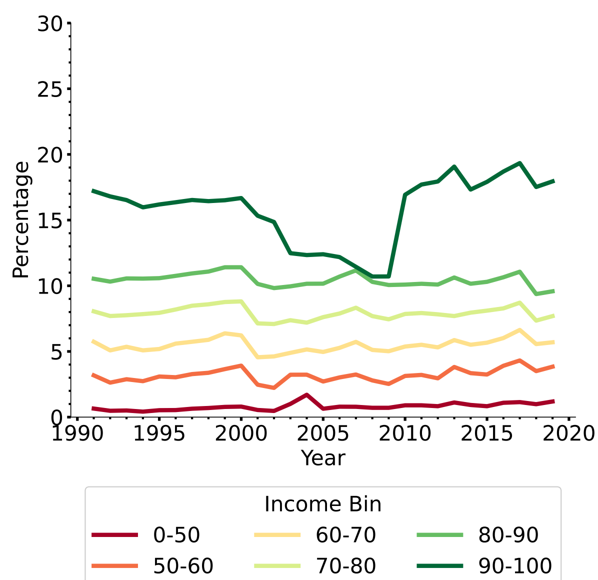

(d) Average Tax Rate on Non-Benefit Income

Notes: This figure shows statistics on taxation by year and by income level. These data are computed on the population of households with Social Security income in the Current Population Survey. The legend gives the percentile bounds of household adjusted gross income for each bin: the bottom half comprise one bin, and the rest are sorted into deciles. Figure A.4a shows the percent of beneficiaries in each bin who are subject to the benefit tax, while Figure A.4b shows the average percent of benefits which are includable in taxable income. Figure A.4c shows the average tax rate on benefits, which is the total benefit tax liability divided by the total level of benefits. Figure A.4d shows the average tax rate on all non-benefit income, similarly defined as the total tax liability on this income divided by the level of income.

Figure A.5: Estimated Coefficients of Dynamic TWFE Model

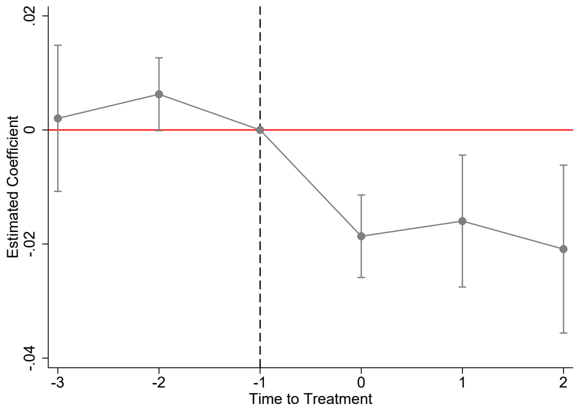

(a) Labor Force Participation

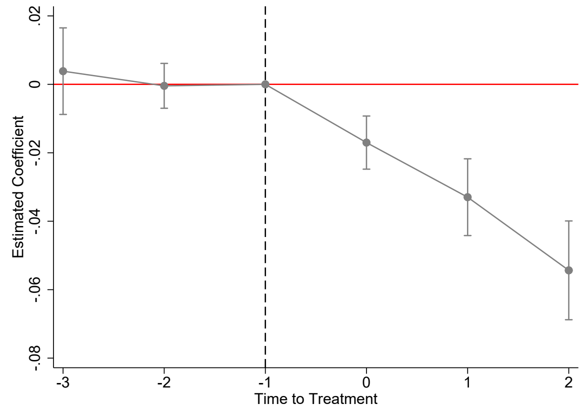

(b) Earnings Conditional on Participation

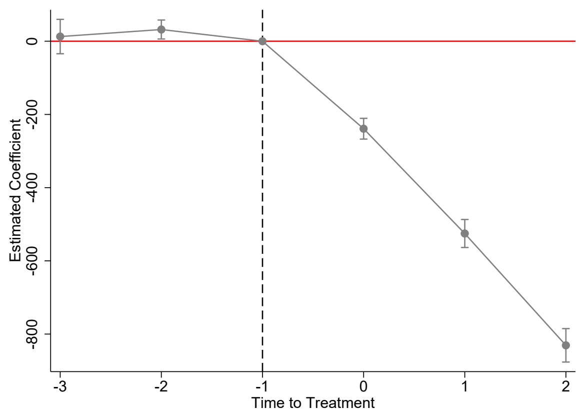

(c) Payroll Taxes

Notes: This figure shows estimated coefficients (and 95% confidence intervals) for the effect of benefit taxation on labor force participation, on earnings conditional on participation, and on payroll tax contributions. Panel A.5a shows the estimated effect of the benefit tax on labor force participation by time to treatment, Panel A.5b shows the estimated effect on the level of earnings for those with non-negative earnings, and Panel A.5c shows the effect on the dollar amount of payroll taxes (including both the employee and employer parts of the tax). An individual is treated when their potential earnings exceed the threshold at which benefits are taxed, and so these results reflect the average effects of any level of taxation. Potential earnings reflect counterfactual earnings without a behavioral response to the benefit tax. Coefficients are estimated using the dynamic TWFE model in Equation 5, see Table A.2 for the exact estimates. Standard errors are clustered at the individual level. This specification includes a control for the natural logarithm of benefits along with fixed effects for individual, time relative to the claim decision, year, sex, and age.

Figure A.6: Annual OASDI Values in SSA and CPS Data

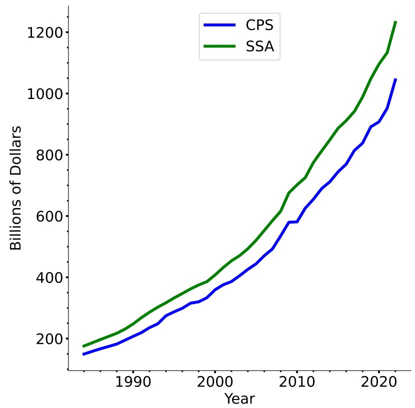

(a) Agg. Benefits, Levels

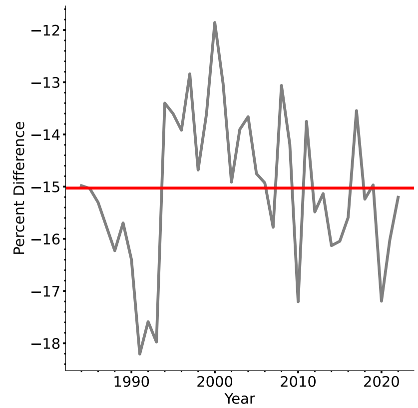

(b) Agg. Benefits, % Diff.

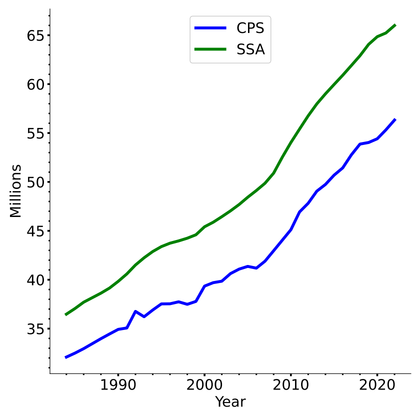

(c) Num. Beneficiaries, Levels

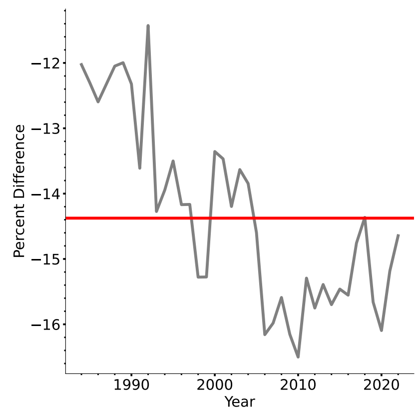

(d) Num. Beneficiaries, % Diff.

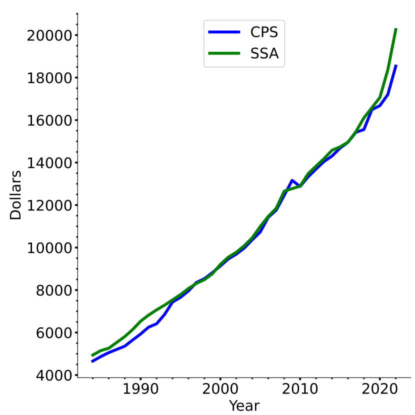

(e) Avg. Benefits, Levels

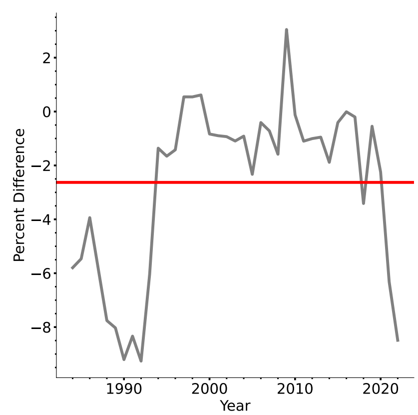

(f) Avg. Benefits, % Diff.

Notes: This figure shows aggregate annual values of amount of benefits, number of beneficiaries, and average benefits per benefits according to SSA summary statistics and according to the CPS-ASEC. Panels A.6a, A.6c, and A.6e show the levels of aggregate benefits, number of beneficiaries, and average benefits, with SSA values in green and CPS in blue. Panels A.6b, A.6d, and A.6f show percent differences between the series, with the red line being the simple average of the percent differences across all years.

Figure A.7: Annual OASDI Values in SSA and CPS Data

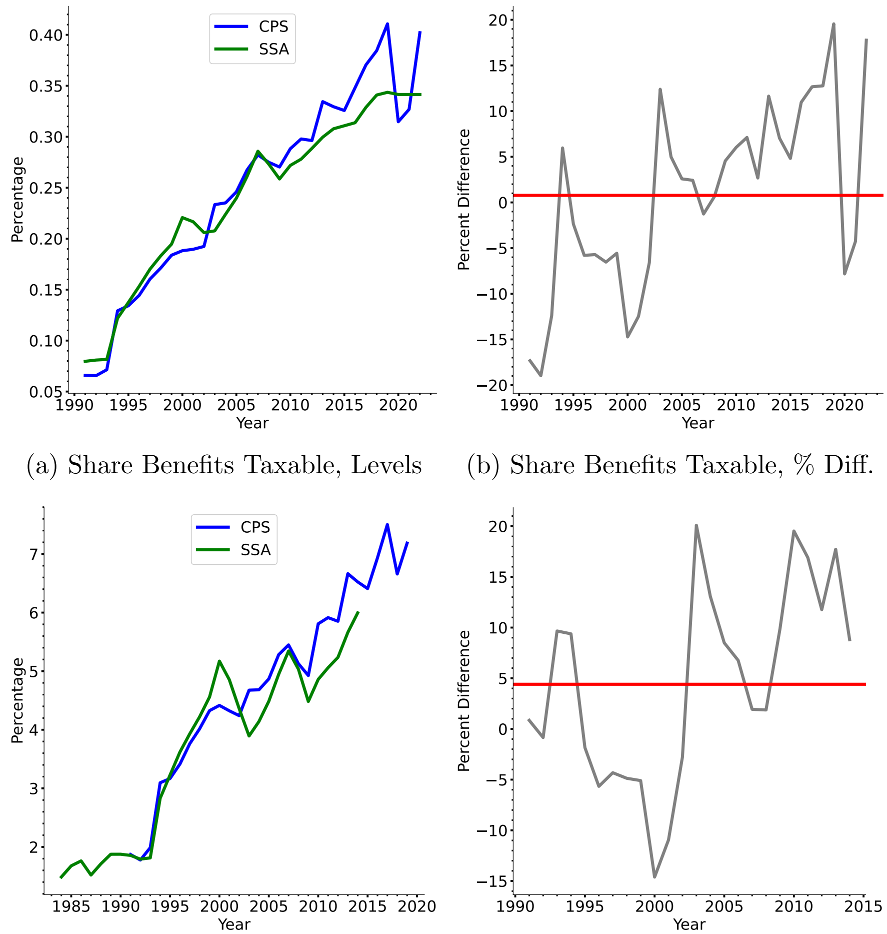

(a) Share Benefits Taxable, Levels

(b) Share Benefits Taxable, % Diff.

(c) Avg. Tax Rate, Levels

(d) Avg. Tax Rate Rate, % Diff.

Notes: This figure shows the closeness of my estimates of aggregate benefit taxation statistics with official data. Panel A.7a shows the percentage of benefits that are includable in taxable income, with my values in blue and SSA's in green. Percent differences between these series are in A.7b. Figure A.7c shows my estimate for the average tax rate on benefits (i.e., total benefit tax liability divided by total benefits) in blue against Social Security in green. Figure A.7d gives the percent differences between these series.

Figure A.8: Estimated Coefficients of Dynamic TWFE Model: Tax Rates

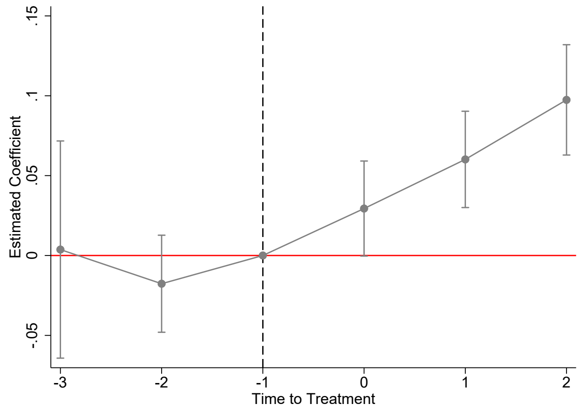

(a) Average Tax Rate

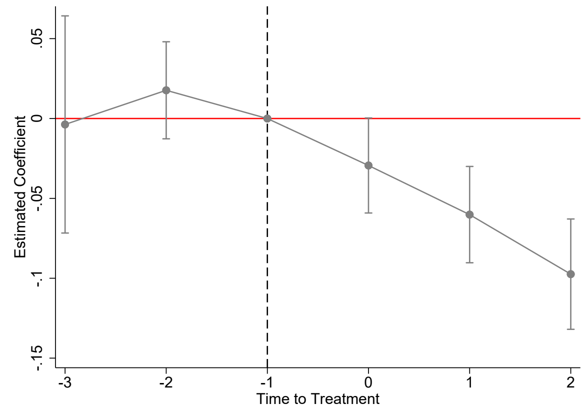

(b) Average Net-of-Tax Rate

Notes: These graphs show estimated coefficients for the effect of benefit taxation on the average tax rate and average net-of-tax rate. The average tax rate is defined as the portion of potential earnings (in the absence of the benefit tax) that are taxed, and the average net-of-tax rate is 1 less this value. An individual is treated when their potential earnings exceed the threshold at which benefits are taxed, and so these results reflect the average effects of any level of taxation. Coefficients are estimated using the dynamic TWFE model in Equation 5, see Table A.4 for the exact estimates. Standard errors are clustered at the individual level. This specification included a control for the natural logarithm of benefits along with fixed effects for individual, time relative to the claim decision, year, sex, and age.

Table A.1: Benefit Tax Thresholds

Threshold	Introduced	Single Phase-In	Joint Phase-In	Rate of Phase-In	Phase-Out
First	1984	$25,000	$32,000	0.5	50% of benefits
Second	1994	$34,000	$44,000	0.85	85% of benefits

Notes: This table shows key parameters that determine the amount of benefits includable in taxable income. This depends on a quantity called modified adjusted gross income (MAGI), which is taxable income plus one-half of Social Security benefits. When MAGI exceeds certain thresholds, benefits are added to taxable income. There are two thresholds, created in 1984 and 1994, which were set in nominal terms and have remained unchanged ever since. The thresholds vary across single and joint filers. For every dollar MAGI exceeds $25,000 ($32,000) for single (joint) filers, $0.50 of benefits are added to taxable income, until 50% of benefits are taxable. Then, for every dollar MAGI exceeds $34,000 ($44,000) for single (joint) filers, $0.85 of benefits are added to taxable income, until 85% of benefits are taxable. The tax rates which apply to these taxable benefits are described in detail below.

Table A.2: Regression Results from Static and Dynamic TWFE Models

	Labor Force Participation		Intensive Earnings		Payroll Taxes	
	Stat.	Dyn.	Stat.	Dyn.	Stat.	Dyn.
Treatment	-0.022		-0.024		-455.410	
	(0.004)		(0.004)		(13.828)	
Treatment - 3		0.002		0.006		13.913
		(0.007)		(0.007)		(24.420)
Treatment - 2		0.007		0.005		73.328
		(0.003)		(0.004)		(13.903)
Treatment + 0		-0.021		-0.019		-308.918
		(0.004)		(0.004)		(15.100)
Treatment + 1		-0.016		-0.033		-610.828
		(0.006)		(0.006)		(19.858)
Treatment + 2		-0.018		-0.052		-841.473
		(0.008)		(0.008)		(24.496)

Notes: This table contains regression results on the effect of benefit taxation on labor force participation, on earnings conditional on participation, and on payroll taxes. Standard errors are given in parentheses and are clustered at the individual level. The payroll tax outcome covers both the employee and employer part of the tax. An individual is treated when their potential earnings exceed the threshold at which benefits are taxed, and so these results reflect the average effects of any level of taxation. Potential earnings reflect counterfactual earnings without a behavioral response to the benefit tax. Both models include the necessary fixed effects for individual and time relative to the Social Security claim decision along with fixed effects for year, sex, age, and a control for the natural logarithm of benefits. The first columns under each outcome variable are estimates under the static model, while the second columns under each outcome allows for the treatment effect to change by time relative to treatment. Negative values indicate periods before treatment, and non-negative values indicate periods after treatment. The period before treatment is omitted.

Table A.3: Elasticities in Response to the Benefit Tax

	With respect to ATR				With respect to ANTR					
	Labor Force Participation		Intensive Earnings		Labor Force Participation		Intensive Earnings		Payroll Taxes	
	Stat.	Dyn.	Stat.	Dyn.	Stat.	Dyn.	Stat.	Dyn.	Stat.	Dyn.
Treatment	-0.069		-0.087		0.330		0.413		-0.593	
Treatment + 0		-0.115		-0.121		0.548		0.576		-0.262
Treatment + 1		-0.041		-0.099		0.193		0.470		-1.068
Treatment + 2		-0.034		-0.117		0.163		0.558		-1.800

Notes: This table shows estimated elasticities of labor supply variables in response to benefit taxation. Elasticities of labor force participation and earnings conditional on participation are presented with respect to the average tax rate and the average net-of-tax rate. The average tax rate is defined as the portion of potential earnings (in the absence of the benefit tax) that are taxed, and the average net-of-tax rate is 1 less this value. For all variables, estimates are presented for the static TWFE model and the dynamic TWFE model. In the dynamic model, treatment effects are given for time of treatment and up to two periods after treatment.

Table A.4: Regression Results from Static and Dynamic TWFE Models: Tax Rates

	ATR		ANTR	
	Stat.	Dyn.	Stat.	Dyn.
Treatment	0.064		-0.064	
	(0.011)		(0.011)	
Treatment - 3		-0.019		0.019
		(0.028)		(0.028)
Treatment - 2		-0.031		0.031
		(0.017)		(0.017)
Treatment + 0		0.037		-0.037
		(0.015)		(0.015)
Treatment + 1		0.077		-0.077
		(0.017)		(0.017)
Treatment + 2		0.103		-0.103
		(0.021)		(0.021)

Notes: This table contains regression results on the effect of benefit taxation on the average tax rate and average net-of-tax rate. The average tax rate is defined as the portion of potential earnings (in the absence of the benefit tax) that are taxed, and the average net-of-tax rate is 1 less this value. Standard errors are given in parentheses and are clustered at the individual level. An individual is treated when their potential earnings exceed the threshold at which benefits are taxed, and so these results reflect the average effects of any level of taxation. Both models include the necessary fixed effects for individual and time relative to the Social Security claim decision along with controls for the natural logarithm of benefits along with year, sex, and age fixed effects. The first columns under each outcome variable are estimates under the static model, while the second column under each outcome allows for the treatment effect to change by time relative to treatment. Negative values indicate periods before treatment, and non-negative values indicate periods after treatment. The period before treatment is omitted.