The Political Economy of Social Security Reform

By Michael J. Boskin, Diego J. Perez, and Daniel S. Bennett*

Appendix

A1. Background Information, Behavioral Assumptions, and Benefit Calculations

Although we focus mainly on the Net Present Value of lifetime and future benefits and taxes, this is just one measure of how people may value the program. Retirement benefits from Social Security are paid out until the death of the beneficiary, i.e. they are indexed annuity payments insuring against the possibility of outliving retirement savings and also against inflation. People may value the program differently than their Net Present Value. Moreover some, perhaps many, workers may not even have accurate information about their prospective benefits. The SSA mails benefit estimates to workers, but has successively been cutting back on who receives mailings to cut costs. As of 2017, only workers over age 60, who do not have an online account with the SSA get them. Other workers under this age must create an online account with the SSA and request the information if they want to see an estimate of their benefits.

The Hoover Institution Survey on Social Security gives us a clearer picture of what people do and do not understand about Social Security. Although 63% understood their benefits were tied to their earnings, 52% believed that benefits rise less than the cost of living, and only 23% knew that the replacement rate is higher for lower income workers. Self help books on maximizing Social Security benefits have been best sellers,[1] suggesting at least some part of the population finds the program difficult to understand. Other surveys, such as Gustman and Steinmeir (2005) and Lusardi and Mitchell (2008), have reported that few respondents seem to understand their Social Security benefits, and in experiments, providing information on Social Security can have measurable impacts on behavior (Liebman and Luttmer, 2011) For these and other reasons, we view the Net Present Value calculations as only a sensible starting point in measuring how people value Social Security (Boskin and Shoven, 1987).

The "deal" that potential voters get from Social Security, both under current law and under the proposed reforms, relies on a number of assumptions including the future rate of real wage growth, future rate of inflation, how people discount future taxes and benefits, and others. We also use simplifications, for example, ignoring complex "file and suspend" strategies some individuals can

* Boskin: Stanford University, Landau Economics Building 579 Serra Mall Stanford, CA, 94305, boskin@stanford.edu. Perez: New York University, Economics Department 19 W 4th Street 6FL New York, NY 10012, diego.perez@nyu.edu. Bennett: Stanford University, Landau Economics Building 579 Serra Mall Stanford, CA, 94305 dsbennet@stanford.edu.
We thank Pablo Villanueva, John Cherian, Daniel Wright, Daniel Begovich, and Neal Begovich for their excellent work as research assistants, participants in the Hoover Institution Economic Policy Working Group Seminar for helpful suggestions, Dave Brady and Doug Rivers for assistance with the Hoover Institution Survey on Social Security, and Daron Acemoglu, Mordecai Kurz, and Roger Myerson for helpful game theory insights.

[1] See Kotlikoff, Moeller, and Solman (2015).

use to maximize benefits. Because of these assumptions and simplifications, the outcomes of the referendums should not be seen as exact predictions of the outcome should such a referendum be held, but as strongly suggestive approximations of how segments of the population will respond to proposed reforms.

All our benefit calculations are based on chapter 7 of the Social Security Handbook, "Figuring the Cash Benefit Rate."[2] Benefits are paid on a monthly basis, with the amount determined by the type of benefit, age when benefits are first claimed and average wages over the worker's lifetime. For retirement benefits, first, the worker's earnings prior to age 60 are "indexed up" by the Average Wage Index (AWI) to the year in which the worker turned 60. Earnings after age 60 are not indexed. The 35 highest earning years are averaged together to get a quantity called Average Indexed Monthly Earnings (AIME). The AIME is used to calculate a baseline Primary Insurance Amount (PIA) by taking 90% of the AIME up to $895 dollars, 32% of the AIME between $895 and $5,397, and 15% of the AIME over $5,397 dollars up to the taxable maximum. The two points at which the slope of the PIA formula changes are called bend points, and they are indexed by the AWI. Finally the PIA is used to determine the monthly benefit amount.

Adjustments are made to the PIA based on the age of the claimant. Every worker has a Full Retirement Age (FRA), determined by the worker's birth year. Historically, the FRA was 65, but the 1983 reform gradually increased it, beginning with individuals turning 62 in 2000, by 2 months per year from 65 to 66. Congress included an 11 year hiatus in the increase, so that people turning 62 between 2005 and 2016 have an FRA of 66, but starting with those turning 62 in 2017, the FRA again increases by 2 months per year until it reaches 67 for those turning 62 in 2022. The earliest a person can claim retirement benefits is age 62; however, in an attempt to make benefit payments actuarially fair regardless of the age at which the worker makes a claim, benefits are permanently reduced by 5/9ths of one percent for each month prior to a workers FRA that the claim is made. Similarly, delaying retirement past a person's FRA increases benefits by 8% per year up to age 70. After applying any early retirement penalties or delayed retirement credits to the worker's PIA, the resulting number is the worker's monthly benefit amount. Monthly payments are thereafter annually adjusted by the change in the CPI-W, attempting to provide beneficiaries with constant real benefits.[3]

Two other types of Social Security benefits are important to our discussion, namely spousal benefits and disability benefits. There are a number of differences in how these benefits are calculated, but we point out just the major differences here. The spouse of an insured worker can claim benefits on their spouse's work history. A person claiming spousal benefits is entitled to one-half of their spouse's benefit while their spouse is alive and their spouse's full benefit after their spouse has died. This holds true for divorced individuals if the marriage lasted for at least 10 years. Married individuals are special also because many of them are entitled to benefits both under their own work history and their spouse's work history. A person eligible for two benefits receives the maximum of the two, not both. Disability benefits are calculated in a

[2]Online available at https://www.ssa.gov/OP_Home/handbook/handbook.html.

[3]The CPI-W likely has overstated inflation by a significant amount cumulatively (Boskin et al. 1996). The CPI research series adjusts the CPI historically to conform to current methodology, a downward adjustment of just under one half of one percent per year. Various chained series are now available that adjust for substitution bias. Finally an experimental CPI-E that attempts to account for the consumption basket of the elderly is also available, although there are several conceptual and measurement problems specific to this index.

similar fashion to retirement benefits. Fewer years are used to calculate the AIME, depending on when the worker becomes disabled, and earnings are indexed to two years prior to the worker's claim, rather than to age 60, when calculating the AIME. Additionally, the monthly benefit is not subject to "early retirement" penalties based on a worker's FRA. When a disabled beneficiary reaches their FRA they are switched from the disability trust fund to the old age trust fund. This is a matter of accounting. The beneficiary receives the same benefits, those benefits are just paid from a different trust fund.

The Disability and Old Age Trust Funds are legally separate entities. The trust funds cannot share revenue or reserves except through Acts of Congress, and the Social Security Administration (SSA) does report on their financial status separately. However, Congress has repeatedly authorized transfers of tax revenues to shore up stability in one trust fund or the other, most recently to prevent the DI Trust Fund from becoming exhausted in 2016, without changing the overall rate of payroll taxes, and much of the analysis by the SSA and others treats them as one combined trust fund. Following most of the SSA's analysis, we ignore the distinction. The Social Security 2100 Act would officially and permanently combine the Old Age and Disability Trust Funds.

After calculating benefits and taxes according to these rules we sum them up to get an expected present value, with the expectation taken over the probability of surviving to receive and pay the benefits and taxes respectively. The interaction of discounting with these survival probabilities leads to the Net of Benefits and Taxes generally being larger for younger archetypes. Consider, for example, the case of a 35 year old and a 45 year old in the medium income cohort. Real wages are growing, so even though they are in the same income cohort, the 35 year olds earn about $8,000 more than the 45 year olds in their highest earning years. But the driving factor of the difference between expected benefits and expected taxes is the different survival probabilities of cohorts. When calculating expected NPV, the expectation is taken given that the cohort has lived to 2019. A person in the 45 year old cohort calculates the probability of living to age 46 given she has lived to age 45. That probability is 0.997, while the probability of an individual in the 35 year old cohort living to 46, given they have lived to age 35, is only 0.977. An individual in the 35 year old cohort has 11 years in which she may die before reaching age 46, while an individual in the 45 year old cohort has only 1 year. This decreases the NPV of taxes of the 35 year olds relative to the 45 year olds because the 35 year olds are less likely to live to pay them. But the probability of living to reach some age is not always lower for the 35 year olds. Younger generations are living longer than previous generations. So while the 35 year olds are about 2 percentage points less likely to live to 46, they are also about one and a half percentage points more likely to live to age 80. The 35 year olds have relatively lower survival probabilities while paying taxes, and relatively higher survival probabilities while collecting benefits, therefore the difference between benefits and taxes tends to be bigger for younger cohorts.

A2. Voting

Voter participation varies along many dimensions in the United States, and some elections depend just as much on convincing supporters to turn out as they depend on convincing undecided voters. To take account of relative voting proclivities, we estimate a probability of turning out to vote for each archetype. Using the November supplement to the Current Population Survey,

in even years from 2008 to 2016, we calculate the share of each archetype[4] that voted in each year, and take the average.[5] We call this American Voting Proclivities, then show that outcomes of the hypothetical referendums on Social Security reform are sensitive to the voting patterns of the population, especially young voters, by considering several alternative voting proclivity assumptions described in the body of the paper.[6]

As discussed in the body, of the paper people vote for representatives rather than individual policies, and a voter may choose to vote for a less preferred Social Security policy because of some other issue that is deemed more important. Despite this reality, occasionally single issues do separately dominate a policy debate and disproportionately affect elections. Also, voters may have more than two policy options to consider. Voting in these situations may be done sequentially, which raises the possibility of strategically voting for less preferred policies. Voting for representatives is conducted geographically by congressional districts, not nationally. Since prices vary markedly across different geographies, the purchasing power of any level of Social Security benefits will also differ. The Commerce Department calculates Regional Price Parities to quantify the difference. For example, with a national average of 100, Cincinnati has an RPP of 89.6 whereas San Francisco stood at 124.7 in 2018. We abstract from these issues, but believe all are interesting opportunities for future research.

We also for the most part assume people vote in their own self interest without consideration of the interests of others. Although we examine how parent's concern for their children's interests could impact referendums, alternatively, we could also consider cases where adult children care about the effect of reforms on their parents. Fourteen percent of adults living in someone else's household are the parents of the household's head, double the 7% rate recorded in 1995 (Kiger 2018). Some younger voters may prefer a more generous Social Security program either out of altruism or to avoid such a situation. More specifically, if the welfare of older generations figures into the reform preferences of their children, our results showing that the preferences of the old tend to dominate referendums would be reinforced.

A3. Determining Income Cohorts for Married, Divorcees, Widow(er)s, and Children

To determine the income cohort of a married couple we use the pair of income cohorts into which each individual falls. We designate a primary earner who can fall into the high, medium or low income cohorts and a secondary earner who can also fall into these cohorts with the additional possibility of having no earnings and the restriction that the secondary earner is never in a higher income cohort than the primary earner.[7] The pairing generates 9 income cohort pairs

[4]The November Supplement contains much less detailed income information. Due to this limitation, we use total household income, which includes some income not subject to the payroll tax, and combines the earnings of two earner households, to sort individuals into income cohorts.

[5]Historically, voter turnout in Presidential elections is about 20 percentage points higher than in midterm elections. Our estimate includes both Presidential and midterm elections. As a sensitivity analysis, we also calculate voting probabilities using only Presidential elections and only midterm elections. Neither changed the outcome of any referendum.

[6]Comparing the November CPS to actual ballot data shows that self reported voting exaggerates voter turnout in the United States. Despite this weakness, the November CPS is still considered to be the best source for comparing the demographic characteristics of those who vote and those who do not. There are a variety of strategic considerations in voting behavior that we ignore for simplicity's sake.

[7]We recognize that individual experiences will vary widely and that over time the role of primary earner may change in many relationships, but the simplification makes the analysis of voting behavior much more tractable.

for married archetypes.

Individuals who have *ever been* married, including divorced, widowed, and separated individuals are assigned to the married cohort, because they are generally still eligible for benefits on their spouse's work history.[8] We assume that married archetypes value their spouse's benefits as much as their own, and so for these archetypes their net present value of the program is equal to the sum of both spouse's net present values. Both spouse's fall into the same age cohort, and retire at the same age. While spouses may be of different ages, a sesitivity analysis assuming spouses are several years apart makes no difference.

Since we cannot observe the spouse's income cohort for separated, divorced, and widowed individuals, we either have to leave them out of the analysis or put them in income cohorts with assumptions that will necessarily misclassify some of them. We choose the latter option. When a divorced individual is observed we assume they are the primary earner and that their spouse does not have zero income, then assign the person randomly among income cohorts that satisfy those two conditions.[9] For individuals too young to have income, we simply use the income cohort of their parents, e.g. a child with a parent in the high income cohort will age into the high income cohort, though in real dollar amounts their AIME will be different. Note that we do not attempt to account for relative regression to the mean. Parental income is an imperfect predictor of an individual's income. For example, Chetty et al. (2014) find that a 10% increase in parental income is associated with only a 3.4% increase in a child's income. A report published in the Federal Reserve Bank of Saint Louis' *Regional Economist*, suggests that the quintile of permanent earnings a child is most likely to end up in, is the same quintile as his/her parent's quintile (Gayle and Hincapé 2016). A more complex analysis might use a full set of transition probabilities, but we deemed that level of detail unnecessary for our purpose. Finally, we compare the distribution among income cohorts to the SSA's public use micro data to ensure that the ultimate distribution of lifetime earnings closely approximates the SSA observations. We make some small adjustments in income cohort cutoffs, especially to the youngest cohorts, to ensure the consistency of the data in that sense.

A4. Archetypes and the Determination of Earnings Histories

We used the method of scaled factors to determine hypothetical work histories for all of our worker-beneficiary archetypes. This is the method used by the SSA, and it is described in their recurring Actuarial Note Number 3. We provide a brief explanation and describe some slight adjustments made for our purposes.

We first note that an alternative method would be to calculate benefits and taxes for a sample of the population rather than for hypothetical workers. That approach would have the advantage of using intra-archetype covariance in the relevant variables, but we would have to split up the sample into something resembling these archetypes anyway to be able to analyze anything other than the total votes. Using archetypes simplifies the calculations without losing much descriptive

[8]Divorced individuals are not eligible for benefits on their spouse's work history if they were married for less than 10 years, or if they remarried. Since we do not observe the lengths of individuals' marriages we make the simplifying assumption here that all individuals in the married cohort are eligible for spousal benefits.

[9]This assumption gave a distribution of income cohorts that most closely resembled the distribution of the SSA's public use microdata.

value.

A scaled factor is the percentage of the AWI a worker earns at a given age. These percentages are calculated from the Continuous Work History Sample, a one percent sample of workers who had some OASDI taxable earnings in their lifetime. For each worker in each year, the SSA calculates the worker's earning as a percent of the AWI. The SSA then takes the average of these percentages by age to obtain, "raw scaled factors." For example, imagine the sample consisted only of two people. Person A born in 1950 and person B born in 1960. Suppose further that person A had earnings equal to 25% of the AWI in 1971 when he was 21, and person B had earnings equal to 75% of the AWI in 1981 when she was 21. The raw scaled factor for age 21 would be the average of 25% and 75%, specifically 50%. Given a year of birth these raw scaled factors imply a level of earnings in each year, which determines a level of average lifetime earnings. Final scaled factors are obtained by multiplying the raw scaled factors by a constant adjustment factor to target a certain level of average lifetime earnings. Notice that this implies that the cutoffs that determine which income cohort an individual falls into depend on the person's age.

We make several small adjustments to the SSA's method. The SSA's scaled factors stop at age 64 when they assume the hypothetical worker retires. We assume people retire probabilistically between ages 62 and 70 in our NPV calculations, so we extend the scaled factors to age 69. Since those who work after age 62 are likely to be very different than the rest of the population, the SSA assumes wages decline geometrically after age 62. We assume the same rate of geometric decrease, and extended it to age 69. We also choose adjustment factors such that workers in our five income groups have average lifetime earnings equal to 40%, 65%, 100%, 145%, or 190% of the AWI in the year the archetype turns 65.

As an alternative to these earnings histories, we also perform our analysis assuming that people earn a constant percent of the AWI. We choose the percentages so that the archetypes have the same average earnings over their lifetime, whether we use scaled factors or not. Table A1 shows the outcome of all pairwise referendums under this assumption. None of the referendums have an outcome different from the base case shown in the body of the paper.

Table A1—: Pairwise Hypothetical Referendums: Assuming Wages are a Constant Percent of the AWI

	Ben18	B18N	Tax18	Price	Ben34	B34N	Tax34
Ben18	-	W[1]	L	L	L	L	L
B18N		-	L	L	W	L	L
Tax18			-	L	W	L	L
Price				-	W	L	L
Ben34					-	L	L
B34N						-	L
Tax34							-

notes: "W" ("L") indicates that the row wins (loses) the hypothetical referendum against the column. Voters participate at the historical levels of American voters.

A5. Disability Insurance

One assumption that deserves special note is the treatment of disability insurance. While this part of the program may not be very salient to most younger workers, its importance has been increasing in recent years. The distinction between taxes dedicated to the Disability Trust Fund, as opposed to the Old Age Trust Fund, may be lost on potential beneficiaries. Pay stubs and W-2 forms do not separately list the taxes due to each trust fund, and the share each trust fund receives is not particularly stable; as mentioned above, in 2016, taxes were reallocated to the disability trust fund to prevent insolvency. Calculating disability benefits introduces a host of complexities including multiple disability spells, incidence by age and gender, and substitution from other benefit programs, and to Medicare. Some States have programs that explicitly try to move people from receiving partly state funded Medicaid, to receiving Federal Disability, and therefore Medicare, which is federally funded. In our baseline analysis we abstract from Medicare and assign to each archetype the actuarial value of their DI benefits, assuming a known constant incidence for all age cohorts and that the disability spell begins at age 55.[10] Here, we show the results under two alternatives here. First, we performed the analysis leaving out disability benefits entirely. Secondly, we split archetypes into those receiving disability benefits and those not receiving disability benefits, calibrating the population shares of each archetype to the incidence of disability estimated by the SSA. The outcomes of the referendums assuming archetypes do not consider disability benefits when choosing a reform to vote for are contained in table A2. None of the outcomes change from the base case.

Table A2—: Pairwise Hypothetical Referendums: Disability Benefits Excluded

	Ben18	B18N	Tax18	Price	Ben34	B34N	Tax34
Ben18	-	W[1]	L	L	L	L	L
B18N		-	L	L	W	L	L
Tax18			-	L	W	L	L
Price				-	W	L	L
Ben34					-	L	L
B34N						-	L
Tax34							-

notes: "W" ("L") indicates that the row wins (loses) the hypothetical referendum against the column. Voters participate at the historical levels of American voters.

Our second approach, which we call deterministic enrollment, was to split archetypes into disabled and non-disabled sub archetypes. Disabled sub archetypes received the NPV and FNPV

[10]Problems analogous to the funding shortfall for Social Security also exists for Medicare and Medicaid. We do not explore those issues here, since a large fraction of these programs are funded by general revenues. In an actuarial sense, most recipients get a good deal given the progressive structure of income tax payments.

of the disability insurance program, while the non disabled sub archetypes received the NPV and FNPV of the old-age insurance program. We calculated the population share of each sub archetype by applying Bayes' Rule to Table 1 of the SSA's "Beneficiary-Characteristics of Disability Insurance" publication, which reports the distribution of age conditional on being a disabled beneficiary. The unconditional distribution of age and disability status were obtained from the Current Population Survey and the "Annual Statistical Report on the Social Security Disability Insurance Program," respectively. Table A3 shows the outcome of the referendums under deterministic enrollment in disability insurance. The outcomes are the same as in the base case.

Table A3—: Pairwise Hypothetical Referendums: Deterministic Enrollment in Disability Insurance

	Ben18	B18N	Tax18	Price	Ben34	B34N	Tax34
Ben18	-	W[1]	L	L	L	L	L
B18N		-	L	L	W	L	L
Tax18			-	L	W	L	L
Price				-	W	L	L
Ben34					-	L	L
B34N						-	L
Tax34							-

notes: "W" ("L") indicates that the row wins (loses) the hypothetical referendum against the column. Voters participate at the historical levels of American voters.

A6. Alternative Macro-economic assumptions

There is a considerable amount of uncertainty regarding the long range values of key macroeconomic variables. The body of the paper uses the SSA's intermediate assumptions unless explicitly stated otherwise, but since these assumptions have a large effect on the size of the actuarial deficit and the archetypes' NPV and FNPV, we perform the analysis using two alternative sets of assumptions. First, we perform the analysis using the CBO's assumptions from their, "Long-Term Budget Outlook." The CBO finds that the actuarial deficit is considerably larger than the SSA does, so to implement their assumptions we not only change the values of key variables, but also the size of the benefit cuts and tax increases. Price indexing of bend points is insufficient to close the deficit under the CBO's assumptions, but its most recent evaluations of potential reform proposals suggests that a 3 year increase in the full retirement age combined with price indexing would approximately close the deficit. The increase in the FRA is phased in by adding two months per year to each individual's FRA, beginning with those first eligible for benefits in 2023, when the most recent increases in the FRA are scheduled to end.

A key driver of the CBO's larger actuarial deficit is the decrease in the share of covered wages below the taxable maximum; in other words, the CBO predicts that wage inequality will increase more than the SSA does. We model an increase in inequality by changing the rate at which real

wages grow for the lowest income and highest income cohorts. We assume a rate of real wage growth of 1% for the lowest income cohort, and calibrate the value for the highest income cohort to maintain an overall rate of real wage growth equal to the CBO's assumed 1.3%. For the highest earners real wages grow by a little more than 1.4% per year. Table A4, shows the results of the hypothetical referendums. Two referendums have outcomes that are different from the base case, namely Ben18 vs Ben31 and Ben31N vs Tax31. Ben18 vs Ben31 switches primarily due to the earlier exhaustion of the trust fund. The earlier exhaustion date makes the delayed benefit cut significantly worse for the 45 and 55 year old cohorts who supported the delayed benefit cut to take advantage of having early years of full benefits. The difference in outcome for Ben31N vs Tax31 is discussed in the body of the paper.

Table A4—: Pairwise Hypothetical Referendums: CBO Long-Term Budget Assumptions

	Ben18	B18N	Tax18	Price	Ben31	B31N	Tax31
Ben18	-	W[1]	L	L	W	L	L
B18N		-	L	L	W	L	L
Tax18			-	L	W	L	L
Price				-	W	L	L
Ben31					-	L	L
B31N						-	W
Tax31							-

notes: "W" ("L") indicates that the row wins (loses) the hypothetical referendum against the column. Voters participate at the historical levels of American voters.

We also performed an analysis of the hypothetical referendums using the high cost scenario assumptions from the SSA's report. The high cost assumptions are intended to indicate a worst case scenario for the trust fund. It does not represent a specific confidence interval, nor is it calibrated to correspond to one, although in simulations the high cost scenario falls outside the 95th percentile of outcomes. We implement the high cost assumptions in the same way as the CBO assumptions. Since the taxable ratio is very similar under the high cost and intermediate cost scenarios, the same real wage growth rate is used for all income cohorts. The results are contained in table A5. We excluded price indexing from the analysis due to the lack of convincing estimates of its effect on the actuarial deficit under these assumptions. The outcomes are the same as under the CBO assumptions.

Another key assumption of our analysis is the rate at which individuals discount future transfers. If the current trend of very low real interest rates on Government Bonds continues, it could very well be the case that our discount rate of 2% is too high.[11] To test the sensitivity of our results to this particular assumption, we perform the analysis again assuming a discount rate of 1%. Tables A6 and A7 show the Lifetime Net Present Value calculations for the base case and

[11]Recall that the discounting also adds mortality probabilities.

Table A5—: Pairwise Hypothetical Referendums: SSA High Cost Scenario

	Ben18	B18N	Tax18	Price	Ben31	B31N	Tax31
Ben18	-	W[1]	L	-	W	L	L
B18N		-	L	-	W	L	L
Tax18			-	-	W	L	L
Price				-	-	-	-
Ben31					-	L	L
B31N						-	W
Tax31							-

notes: "W" ("L") indicates that the row wins (loses) the hypothetical referendum against the column. Voters participate at the historical levels of American voters.

the case with a low discount rate respectively. The discount rate makes a significant difference in the level of LNPV. For example, a single, 45 year old in the low income cohort has a present value of benefits equal to about 3.8 times his income at age 65 when using a 2% discount rate, but 5.1 times his income when using a 1% discount rate. For a single high income 45 year old, the benefits go from being about 2 times her income to 2.7 times her income. Importantly, a number of high earning archetypes have negative LNPV in the base case, but a positive LNPV when using the lower discount rate. Despite these significant differences in the levels of LNPV, none of the hypothetical referendums have a different outcome when using the lower discount rate.

A7. Income Taxation of Benefits

The 1983 amendment to the Social Security Act provided for the income taxation of benefits. The law specified that single individuals earning over $25,000 and married couples earning over $32,000 would be required to treat up to half of their benefits as taxable income. A subsequent amendment added two additional thresholds. For individuals and couples earning over $34,000 and $44,000 respectively, up to 85% of their benefits are treated as taxable income. The income taxation of benefits is in some ways a small part of the program. Revenue from the income taxation of benefits accounts for just 3.8 percent of total trust fund revenues. On the other hand, because the thresholds are not indexed to inflation or wages the share of beneficiaries required to pay income taxes on their benefits has been steadily rising. In 1984 10% of beneficiaries owed income taxes on their benefits, while today that number is 56%.

Including income taxation in our analysis requires additional assumptions. Who owes taxes, the average tax rate, and the marginal tax rate all depend on retirement income that is notoriously under reported in the Current Population Survey (Bee and Mitchell 2017). To determine who owes Social Security taxes we make the simplifying assumption that the highest lifetime income cohorts pay taxes on benefits, and we calibrate "pseudo" cutoffs in the model to match the observed and predicted shares of the population paying income taxes on benefits. To determine

Table A6—: Lifetime Net Present Value of Transfers, Discounted to Age 65, for Married Archetypes Assuming no Action on Unfunded Liabilities.

$2019 Thousands

		Income Cohorts								
	Pr	Low	Low	Med	Med	Med	High	High	High	High
	Sp	$0	Low	$0	Low	Med	$0	Low	Med	High
25	Ben	515.4	643.3	898.6	942.9	1121.6	1287.9	1287.9	1425.8	1607.8
25	Tax	224.0	448.1	560.1	784.1	1120.2	1064.2	1288.2	1624.3	2128.4
25	Net	291.3	195.3	338.5	158.8	1.4	223.7	-0.3	-198.5	-520.6
45	Ben	391.2	486.6	682.1	715.2	848.4	977.6	977.6	1080.4	1216.1
45	Tax	178.7	357.4	446.7	625.4	893.4	848.7	1027.4	1295.4	1697.4
45	Net	212.5	129.2	235.4	89.8	-45.0	128.9	-49.8	-215.0	-481.4
65	Ben	319.0	395.0	563.5	587.1	697.6	794.6	794.8	878.9	984.2
65	Tax	142.2	284.3	355.4	497.6	710.8	675.3	817.5	1030.7	1350.6
65	Net	176.8	110.7	208.1	89.5	-13.2	119.3	-22.7	-151.8	-366.4

note: Age categories are 10 year intervals centered on the indicated age. Pr and Sp indicate the income of the primary and secondary earners respectively.

the taxes paid we follow the methods of the Social Security Administration.[12]. Specifically we use estimates from the Treasury Department's Office of Tax Analysis to compute a Ratio of Taxable Benefits, and the average share of individual benefits paid back to the trust fund in income taxes. We treat this as the average tax rate the individuals face on their benefits.

We calculate the expected present value of income taxes for each archetype and add it to the expected present value of payroll taxes. Social Security looks slightly less generous when considering these taxes. Tables A8 and A9 show the FNPV calculations for selected married archetypes. All of the 25 year old archetypes will owe income taxes on their benefits. A comparison of the magnitudes shows that the difference in the future present value of taxes is small even for the highest income cohorts. Importantly, no archetype with a positive FNPV has a negative FNPV when income taxes on benefits are included. Unsurprisingly given the small magnitudes, the conclusions we draw are unchanged. None of the hypothetical referendums has a different outcome compared to the base case where the income taxation of benefits is not considered.

*

REFERENCES

[12] "Long Range OASDI Projection Methodology," 2018

Table A7—: LNPV for Selected Married Archetypes Discounted to Age 65 with a Discount Rate of 1%.

$2019 Thousands

		Income Cohorts								
Age	Pr Sp	Low $0	Low Low	Med $0	Med Low	Med Med	High $0	High Low	High Med	High High
25	Ben	576.9	716.7	1005.9	1054.2	1249.6	1441.7	1441.7	1592.2	1791.2
25	Tax	181.2	362.4	453.0	634.3	906.1	860.8	1042.0	1313.8	1721.6
25	Net	395.7	354.3	552.9	419.9	343.5	580.9	399.7	278.4	69.6
45	Ben	436.0	539.9	760.3	796.2	941.2	1089.7	1089.7	1201.5	1349.2
45	Tax	144.1	288.3	360.4	504.5	720.8	684.7	828.9	1045.1	1369.4
45	Net	291.9	251.6	400.0	291.7	220.5	405.0	260.8	156.4	-20.2
65	Ben	352.8	434.9	623.3	648.7	768.1	878.9	879.2	970.0	1083.6
65	Tax	114.8	229.6	287.0	401.8	573.9	545.2	660.0	832.2	1090.5
65	Net	238.0	205.4	336.3	247.0	194.1	333.7	219.1	137.8	-6.9

note: Age categories are 10 year intervals centered on the indicated age. Pr and Sp indicate the income of the primary and secondary earners respectively.

Table A8—: FNPV for Selected Married Archetypes Discounted to 2019 without Income Taxes on Benefits

$2019 Thousands

		Income Cohorts								
Age	Pr Sp	Low $0	Low Low	Med $0	Med Low	Med Med	High $0	High Low	High Med	High High
25	Ben	515.4	643.3	898.6	942.9	1121.6	1287.9	1287.9	1425.8	1607.8
25	Tax	214.7	429.4	536.8	751.5	1073.6	1019.9	1234.6	1556.7	2039.8
25	Net	300.6	213.9	361.8	191.4	48.1	268.0	53.3	-130.9	-432.0
45	Ben	391.2	486.6	682.1	715.2	848.4	977.6	977.6	1080.4	1216.1
45	Tax	79.4	158.7	198.4	277.8	396.8	377.0	456.4	575.4	754.0
45	Net	311.8	327.9	483.7	437.4	451.5	600.6	521.2	505.0	462.1
65	Ben	289.0	357.9	510.3	531.3	631.8	719.9	720.1	796.0	891.6
65	Tax	1.1	2.3	2.8	4.0	5.7	5.4	6.6	8.3	10.8
65	Net	287.8	355.6	507.4	527.4	626.1	714.5	713.5	787.7	880.8

note: Age categories are 10 year intervals centered on the indicated age. Pr and Sp indicate the income of the primary and secondary earners respectively.

Table A9—: FNPV for Selected Married Archetypes Discounted to 2019 with Income Taxes on Benefits.

$2019 Thousands

		Income Cohorts								
Age	Pr Sp	Low $0	Low Low	Med $0	Med Low	Med Med	High $0	High Low	High Med	High High
25	Ben	515.4	643.3	898.6	942.9	1121.6	1287.9	1287.9	1425.8	1607.8
25	Tax	214.7	429.4	536.8	751.5	1073.6	1019.9	1234.6	1556.7	2039.8
25	Net	300.6	213.9	361.8	191.4	48.1	268.0	53.3	-130.9	-432.0
45	Ben	391.2	486.6	682.1	715.2	848.4	977.6	977.6	1080.4	1216.1
45	Tax	79.4	158.7	198.4	277.8	396.8	377.0	456.4	575.4	754.0
45	Net	311.8	327.9	483.7	437.4	451.5	600.6	521.2	505.0	462.1
65	Ben	289.0	357.9	510.3	531.3	631.8	719.9	720.1	796.0	891.6
65	Tax	1.1	2.3	2.8	4.0	5.7	5.4	6.6	8.3	10.8
65	Net	287.8	355.6	507.4	527.4	626.1	714.5	713.5	787.7	880.8

note: Age categories are 10 year intervals centered on the indicated age. Pr and Sp indicate the income of the primary and secondary earners respectively.