Article

Public Finance Review
1-37

DOI: 10.1177/1091142118770199
journals.sagepub.com/home/pfr
SAGE

Social Security and Saving: An Update

Sita Slavov[1], Devon Gorry[2], Aspen Gorry[2], and Frank N. Caliendo[3]

Abstract

Typical neoclassical life-cycle models predict that Social Security has a large and negative effect on private savings. We review this theoretical literature by constructing a model where individuals face uninsurable longevity risk and differ by wage earnings, while Social Security provides benefits as a life annuity with higher replacement rates for the poor. We use the model to generate numerical examples that confirm the standard result. Using several benefit and tax changes from the 1970s and 1980s as natural experiments, we investigate the empirical relationship between Social Security and private savings and find little evidence to support the predictions from the theoretical model. We explore possible reasons for the lack of strong empirical findings.

Keywords

Social Security, saving, life insurance

[1] George Mason University, Arlington, VA, USA
[2] Clemson University, Clemson, SC, USA
[3] Utah State University, Logan, UT, USA

Corresponding Author:
Sita Slavov, George Mason University, 3351 Fairfax Drive, MS 3B1, Arlington, VA 22203, USA.
Email: sslavov@gmu.edu

Projections by the Social Security Administration suggest that the Old Age and Survivors Insurance trust fund will be depleted in 2035. To shore up the program's finances, some policy makers propose cutting benefits relative to current law. However, other policy makers express concern that benefit cuts will harm lower-income retirees and instead propose expanding the Social Security system by increasing benefits for large segments of the population, with these increases funded by raising payroll taxes. Because Social Security requires contributions from workers and provides benefits during retirement, theory predicts that expanding the program will crowd out private savings for individuals who are not liquidity constrained. Besides retirement income, Social Security also provides life insurance in the form of benefits paid to the dependent children of deceased workers. Theory predicts that expanding Social Security will also crowd out private life insurance purchases among those who are not liquidity constrained. In this article, we review the existing theory and then explore the empirical evidence for its predictions.

Understanding the effect of Social Security on private savings and life insurance holdings is important for a few reasons. First, behavioral responses are an important factor in evaluating how well Social Security meets its objectives. Any potential welfare gains from mandatory saving and risk sharing through Social Security depend on the degree to which households attempt to unwind their Social Security contributions through reductions in private savings and insurance holdings. For instance, Hosseini (2015) shows that the large welfare gains from mandatory annuitization are mostly washed out by the distortions that it causes to the equilibrium price of annuity contracts, as adverse selection causes high mortality individuals to exit the private annuity market. Second, the impact of changes in Social Security rules to maintain fiscal solvency in the face of a demographic shock depends on how individuals respond (Kitao 2014; Bagchi 2016). Finally, if Social Security causes large reductions in personal savings, then growth in the generosity of the Social Security system in recent decades may help to explain the precipitous decline in the aggregate personal saving rate in the United States from near 10 percent in the 1980s to close to zero in more recent years (Parker 1999; Gokhale, Kotlikoff, and Sabelhaus 1996).

To motivate our empirical work, we provide a brief review of the predictions of a basic, neoclassical life-cycle model. We use a variant of Yaari's (1965) classic model to review the theoretical effect of Social Security on private savings by income group. The model features rational individuals who face uninsurable longevity risk and differ by their level of earnings. Social Security provides longevity insurance because it pays benefits as a life annuity. Moreover, it redistributes wealth from high- to low-wage earners because the benefit-earning rule is progressive. We document that Social Security has a large effect on private asset holdings at retirement among all income groups below the tax cap.

Next, we empirically examine the impact of plausibly exogenous variation in Social Security retirement benefits and taxes created by the 1977 and 1983 Social Security reforms as well as a plausibly exogenous reduction in the life insurance value of Social Security that was phased in between August 1981 and April 1985. The 1977 reform reduced benefits for individuals born in 1917 and later, but not for individuals born in 1916 and earlier. The 1983 reform increased payroll taxes for self-employed individuals relative to wage earners; it also increased the full retirement age (FRA) for individuals born in 1938 and later. An increase in the FRA is equivalent to a cut in Social Security wealth. The reduction in the life insurance value of Social Security comes from a 1981 reform that eliminated dependent benefits for college students aged 18 to 21 beginning in 1982 with a phase-out period completed in 1985.

There is a large macroeconomic literature that assesses the impact of Social Security (and Social Security reform) on private savings. This literature uses calibrated, dynamic general equilibrium models (e.g., Blau 2016; Conesa and Garriga 2008; Kitao 2014). Relative to the macroeconomic literature that focuses primarily on aggregate saving behavior, our review of theoretical predictions focuses on the differential impact of Social Security on private savings across different income groups.

Numerous microeconometric studies have also estimated the relationship between Social Security wealth (or pension wealth more generally) and private savings. Some early studies use aggregate time series data (e.g., Feldstein 1974, 1996; Leimer and Lesnoy 1982), while others use micro data (Feldstein and Pellechio 1979; Diamond and Hausman 1984; Bernheim 1987; Bernheim and Levin 1989; Gale 1998; Gustman and Steinmeier 1998; Kotlikoff 1979). These studies find mixed results. Some find that Social Security wealth crowds out private savings almost one for one; others suggest that Social Security wealth reduces private savings somewhat, but not one for one, so that total household savings (public plus private) increases in the presence of Social Security.

One shortcoming of this approach—which relies on examining the correlation between Social Security wealth and private savings, after controlling for observables—is that there are likely numerous unobservable factors that influence both Social Security wealth and saving preferences. Our contribution is to use plausibly exogenous policy-induced variation in

Social Security wealth to estimate the impact on private savings. Similar studies done for other countries have found that public pensions tend to crowd out private savings in Italy, Mexico, China, and the United Kingdom (e.g., Attanasio and Brugiavini 2003; Aguila 2011; Feng, He, and Sato 2011; Attanasio and Rohwedder 2003). However, evidence for the United States is much more limited. Engelhardt and Kumar (2011) examine the impact of defined benefit pension wealth on nonpension savings for the United States, using variation in the benefit accumulation rules for such plans. They find that increases in pension wealth reduce nonpension savings. However, they do not focus on Social Security.

In contrast to previous microeconometric studies based on policy-induced variation in public and private pension wealth, we find little evidence that policy-induced variation in Social Security has affected private savings. There is some evidence that payroll tax increases in 1983 may have reduced savings, but there is not much evidence that the benefit changes enacted in 1977 and 1983 had an impact on savings. We also find no evidence that the elimination of dependent benefits for college students increased life insurance holdings.

Even though our findings are largely null results, this study makes several contributions to the literature. First, it highlights an inconsistency between the strong predictions of theory and the weak empirical results. We provide a detailed discussion of why this inconsistency may arise. Second, the publication of valid, null results can highlight the extent of uncertainty of our knowledge and reduce publication bias (see, e.g., Franco, Malhotra, and Simonovits 2014). Our null results are valid because the methods and data we have used are arguably the best available for the period. In contrast to previous studies for the United States, we exploit variation in Social Security wealth that is plausibly exogenous. In the context of the previous literature, which has generally found stronger effects, our results highlight the uncertainty of our knowledge regarding the impact of Social Security on private savings.

There are several possible explanations for our lack of a finding. First, our standard errors are large, so we cannot rule out the possibility that there is an economically meaningful effect. Second, it is possible that individuals chose to delay retirement, rather than save more, in response to the benefit cuts that we study. Third, it is possible that individuals simply lack knowledge about their Social Security benefits and were unaware of the policy changes. Fourth, the timing of information about the policy changes could play a role in determining when behavioral changes are observed in the data. If the policy changes were anticipated in advance, behavior may have adjusted well before the policy changes were implemented. Fifth, expectations may also affect the timing of behavioral responses. For example, Ricardian equivalence suggests that individuals fully anticipate policy changes designed to restore actuarial balance to the Social Security program. Thus, adjustment to these policy changes may occur long before they are implemented or even discussed publicly. Finally, individuals may not respond to changes in Social Security wealth due to either liquidity constraints or participation in means-tested welfare programs.

The remainder of this article is organized as follows. The second section reviews the predictions from a standard neoclassical life-cycle model. The third section presents our empirical methods. The fourth section presents our empirical results. The fifth section provides a discussion exploring why the empirical findings differ from the theoretical results. The sixth section concludes.

Brief Review of Predictions from a Standard, Neoclassical Life-cycle Model

This section reviews a standard version of a life-cycle consumption savings model to understand the theoretical implications of Social Security on private savings. Before we begin, we emphasize that there is no such thing as *the* neoclassical life-cycle model; instead, it is a framework for analyzing the effects of public policies on decision-making and welfare. The framework assumes that individuals behave rationally given the risks and policies that they face. Beyond this assumption, we do not attempt to include every feature that has been considered in the large literature using such models to study consumption and saving behavior. We focus on a stylized setting where private insurance is missing, capital markets are complete, and labor supply is inelastic.[1]

As in Yaari (1965) and Bütler (2001), individuals in our model face survival uncertainty up to a maximum possible age beyond which survival is impossible. While a simple two-period model could be used to make the same basic points that we will make in this section, our continuous time setting gives us the ability to model realistic survival probabilities, which in turn is important in the study of the life annuity aspect of Social Security. In addition, individuals in our model save only for their own retirement rather than for their descendants. Therefore, our finite-horizon model without a bequest motive contrasts with an infinite-horizon (dynastic) model, or equivalently, a finite-horizon model with a "perfect" bequest motive. While the effect of Social Security on private saving may differ somewhat across these different model settings, we focus on a model in the Yaari and Bütler tradition.

Age is continuous and is indexed by t. At each moment in time, an infinitely divisible cohort of unit mass is born. Individuals are born at $t = 0$ and die no later than $t = 1$. The probability of surviving to age t from the perspective of age 0 is $S(t)$, where $S(0) = 1$ and $S(1) = 0$. Individuals receive an exogenous and constant flow of wage income w up to the exogenous retirement date t_R, and they receive constant Social Security benefits $b(w)$ in the form of a life annuity after retirement. Wage income is taxed at the Social Security tax rate τ. Individuals differ only by wage income. Each individual draws their income from the p.d.f. $g(w)$ with support $[0, 1]$. Including wage heterogeneity in the model allows us to study how a progressive Social Security system affects saving rates throughout the wage distribution.

An individual's consumption is $c(t)$ and flow utility from consumption is $u(c(t))$, where $u' > 0$ and $u'' < 0$. Annuity markets are completely closed, and saving is done in a zero-interest storage technology $k(t)$. The assets of the deceased are bequeathed to the new generation. Bequest income per newborn B is collected at $t = 0$. Factor prices are fixed because we focus on a small, open economy.[2]

We study stationary equilibria in which individuals behave rationally in an environment where they face longevity risk, the Social Security budget is balanced, and the transmission of wealth across generations is consistent in the sense that the aggregate assets of the deceased equal the inheritances of the living. We will compare equilibrium assets at retirement with and without Social Security.

Equilibrium

At $t = 0$, individuals learn their wage type w and take as given τ, $b(w)$, and B. They choose $(c(t), k(t))_{t\in[0,1]}$ according to:

$$\max \int_0^1 S(t)u\Big(c(t)\Big)dt, \tag{1}$$

subject to:

$$\dot{k}(t) = (1 - \tau)w - c(t), \text{ for } t \in [0, t_R], \tag{2}$$

$$\dot{k}(t) = b(w) - c(t), \text{ for } t \in [t_R, 1], \tag{3}$$

$$k(0) = B,\ k(1) = 0. \tag{4}$$

The Social Security system runs a balanced budget and hence the tax rate τ must satisfy the government's budget constraint for a given benefit-earning rule $b(w)$:

$$\tau R \int_0^1 g(w) w dw = \int_0^1 g(w) b(w) dw, \tag{5}$$

where R is the ratio of workers to retirees, $R \equiv \int_0^{t_R} S(t)dt / \int_{t_R}^1 S(t)dt$. There are no inefficiencies in financing Social Security. The government can store wealth at zero interest just like the private market. However, Social Security pays an above-market implicit rate of return because it pools the contributions of the deceased to pay an annuity to the living.

The inheritances of the living must equal the assets of the deceased. We follow the standard assumption that bequest income is spread evenly across the new generation. Equilibrium bequest income B solves the following implicit function:

$$B = \int_0^1 \int_0^1 g(w) \Big(- \dot{S}(t) \Big) k(t|w, B) dt dw. \tag{6}$$

For a given income distribution $g(w)$ and Social Security benefit rule $b(w)$, a stationary equilibrium consists of (i) individual consumption and saving decisions $\{c^*(t|w), k^*(t|w)\}$ for each wage type w that satisfies individuals' optimization problem, (ii) a Social Security tax rate τ^* that balances the government's aggregate budget constraint, and (iii) bequest income B^* that balances the inheritances of newborns with the assets of the deceased.

Numerical Example

The parameters that we need to select to generate numerical examples are the retirement age t_R, the survival function $S(t)$, which in turn pins down the ratio of workers to retirees R, the utility function $u(c)$, the wage density $g(w)$, and the Social Security benefit rule $b(w)$. Once these parameters are chosen, we can solve for the equilibrium quantities $\{c^*(t|w), k^*(t|w), \tau^*, B^*\}$.

We imagine an individual who works from age 25 to 65 and dies no later than age 100. Hence, $t_R = 40/75$. We follow Caliendo, Gorry, and Slavov (2016) in setting the unconditional survival probabilities to $S(t) = 1 - t^{3.28}$.[3] We use constant relative risk aversion (CRRA) utility $u(c) = c^{1-\sigma}/(1-\sigma)$ with $\sigma = 3$.

The Social Security benefit-earning rule is a piecewise linear function of an individual's wage. There are three kinks or bend points which change each year based on average wage growth, but we follow Alonso-Ortiz (2014) and others and assume that the bend points are the following constant multiples of the economy-wide average wage, 0.2, 1.24, and 2.47. Social Security replaces 90 percent of the individual's wage up to the first bend point, 32 percent of wage income between the first and second bend points, and 15 percent of wage income between the second and third bend points. The third bend point, 2.47, corresponds to the maximum taxable earnings (beyond which individuals do not pay payroll taxes or receive additional benefits). Therefore, Social Security replaces 0 percent of wages beyond the third bend point.

Finally, wages follow a beta distribution with density:

$$g(w) = \frac{w^{\gamma-1}(1-w)^{\beta-1}}{\int_0^1 w^{\gamma-1}(1-w)^{\beta-1}dw}, \text{ for } w \in [0, 1]. \tag{7}$$

There are two parameters to choose: γ and β. To abstract from the portion of the income distribution above the third bend point (or maximum taxable earnings), we assume that the top wage earner with $w = 1$ earns 2.47 times that of the average earner. In addition, the 2015 US Census reports a Gini coefficient of 0.479 on the distribution of income. This gives us two targets with which to calibrate γ and β:

$$E(w) = \frac{\gamma}{\gamma + \beta} = \frac{1}{2.47}, \tag{8}$$

$$\text{Gini} = \frac{1}{2E(w)} \int_0^1 \int_0^1 g(x)g(y)|x-y|dxdy = 0.479. \tag{9}$$

By setting $\gamma = \beta/1.47$, we match the desired mean. Given this restriction, $\beta = 1.8$ delivers a Gini coefficient of 0.3463, somewhat less than the desired target but about the best we can do given the (truncated) beta distribution that we have assumed. We would expect the model to

Table 1. Assets at Retirement as a Fraction of Lifetime Income by Wage Type.

	Wage type			
	$0.2E(w)$	$E(w)$	$1.24E(w)$	$2.47E(w)$
No Social Security	41 percent	41 percent	41 percent	41 percent
Social Security	−6 percent	11 percent	12 percent	18 percent

understate the true degree of inequality given that we ignore individuals above the maximum taxable income.

Using this benefit-earning rule, together with the survival function $S(t)$ and the wage density $g(w)$, we find that the balanced budget tax rate is $\tau^* = 19.463\%$. This is larger than the current tax in the United States (10.6 percent), but the current rate is not sufficient to balance the expected future budget under current life expectancies. According to the 2016 Social Security Trustees report, eliminating the infinite horizon actuarial shortfall in the retirement and disability programs combined would require a 4.2 percentage point increase in the payroll tax. That increase—which overestimates the additional payroll tax needed to eliminate the shortfall in the retirement program alone—would still result in a tax rate that is below the one in our model. However, our analysis also ignores the mass of individuals who max out their Social Security contributions each year and would increase the overall revenue. To test for sensitivity, we compute results by setting the tax to 10.6 percent and ignoring the budget balance and find that the qualitative results that we emphasize do not change.

We compare two equilibria. The first equilibrium features a Social Security program as described above. The second equilibrium has no Social Security with $\tau = b(w) = 0$ for all w. Everything else is the same across the two equilibria.

Table 1 reports equilibrium assets at retirement as a fraction of lifetime income (wage income plus bequest income) for individuals at each of the three bend points as well as for average earners. Without Social Security, all individuals save about 41 percent of their lifetime income. The impact of Social Security on private savings is enormous. The effect is especially large among the poor who face the highest replacement rates. Social Security wipes out as much as half of the private savings of those near the earnings cap, and it wipes out all of the private savings of those at the first bend point. Note that while the declines in private savings as a percent of lifetime income are especially large for the poor, the aggregate reduction in dollars saved still may be smaller as their wages are much lower than other income groups considered. Including other features will alter the precise quantitative effect of Social Security on private savings, but if individuals have perfect foresight, they tend to unwind a significant portion of mandatory saving with reductions in private saving. Indeed, if we were to abstract from longevity risk and wage heterogeneity, then Social Security taxation would crowd out private saving one for one.

In addition to private savings, Social Security may affect insurance holdings. In theory, Social Security will reduce the demand for both life insurance and annuity insurance. In our model, bequests are accidental rather than the result of an explicit bequest motive. With the addition of an explicit bequest motive, individuals would demand private life insurance. The magnitude of the demand for life insurance would depend on the strength of the bequest motive and on the level of life insurance provided publicly through Social Security. Hence, Social Security may crowd out private life insurance holdings in addition to its crowding out of private savings (Hong and Ríos-Rull 2007; Li 2016). Similarly, Social Security also reduces the demand for private annuities because the program offers significant annuitization through its payment of benefits as a life annuity. This effect is particularly large for those with high mortality. Since Social Security disproportionately causes those with the highest mortality to exit the private annuity market, it worsens adverse selection problems in the private annuity market (Hosseini 2015).

Empirical Analysis

Empirically testing the theoretical predictions from the previous section requires plausibly exogenous changes to Social Security benefits or taxes. The most recent major changes to the Social Security retirement program were made in 1977 and 1983. A smaller change to dependents' benefits was made in 1981. Since there have been no significant policy experiments since then, we rely on the changes made in the late 1970s and early 1980s to provide identifying variation in Social Security benefits and taxes.[4] While the reforms we study were discussed publicly well in advance, our identification strategy relies on comparing groups that were affected by each reform with groups that were unaffected. Thus, to the extent that the exact design of the reforms—for example, the age cutoffs determining which groups were grandfathered—could not be anticipated, our identifying variation is plausibly exogenous. Another shortcoming of these policy experiments is that none of them represent balanced budget expansions of Social Security of the sort considered in the theoretical

Table 2. Availability of Data and Wealth Measures.

	Years available	Cohorts available
Consumer expenditure survey	1972–1973, 1980–	1890–1976[a]
Panel study of income dynamics	1968–1972, 1979, 1980 (1984, 1989, 1994, and 1999–)[b]	1889–1955[a]
National longitudinal survey	1966, 1969, 1971, 1976, 1981, and 1990	1907–1921[c]
Retirement history survey	1969–1979	1896–1911[d]

[a]Tabulated from data for sample years.
[b]Wealth variables change for 1984 onward.
[c]Aged 45 to 59 in 1966.
[d]Aged 58 to 63 in 1969.

model. All were designed to reduce long-run financial shortfalls. We revisit both issues after presenting the results. Despite these issues, we believe that exploiting policy-induced variation in Social Security wealth is an improvement over examining aggregate time series or individual-level correlations between Social Security wealth and private savings, which is the approach that most of the existing literature has taken.

Unfortunately, data sets like the Health and Retirement Study (HRS) and the Survey of Consumer Finances (SCF)—which are often used today to study saving and retirement behavior—are not available for periods that we study. However, several other micro data sets are available. Table 2 provides information on the only four micro data sets that—to our knowledge—provide information on household-level wealth holdings and saving behavior for the periods studied. These include the Panel Study of Income Dynamics (PSID), the National Longitudinal Survey (NLS), the Consumer Expenditure Surveys (CE), and the Retirement History Survey (RHS). These surveys contain poorer-quality wealth data than the SCF and HRS. However, they are the best available for the period. The theoretical model makes predictions about total savings. Thus, for each survey, we attempt to construct the most comprehensive measures of savings possible.[5]

Our first policy experiment, the 1977 reform, resulted in reduced benefits for individuals born in 1917 and later. The purpose of the reform was to correct a mistake in a 1972 law that had introduced automatic cost-of-living adjustments for Social Security benefits. The indexation formula in the

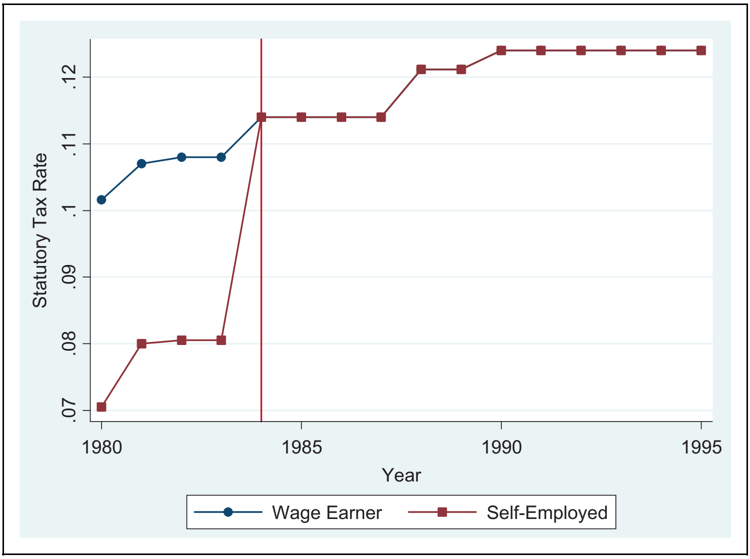

Figure 1. Statutory payroll tax rate by employment status.

1972 law resulted in cost-of-living adjustments that far exceeded inflation for individuals who had not yet claimed. The correction made in 1977 was phased in for individuals born between 1917 and 1921. The lower benefits paid to those born between 1917 and 1921, relative to those born in 1916 and earlier, are sometimes referred to as the "notch," with individuals in the 1917 to 1921 birth cohorts referred to as "notch babies."[6] According to table 2, the data sets that could potentially be used to study the impact of this reform include the PSID and the NLS. We use both in our analysis. The RHS also covers the late 1970s; however, since cohorts born after 1911 are not included in the survey, we cannot use this data set.

Our second policy experiment, the 1983 reform, increased payroll taxes for self-employed individuals relative to wage earners. Figure 1 shows the total statutory Old Age, Survivor, and Disability Insurance (OASDI) payroll tax rate faced by employees (including employer and the employee shares) as well as the self-employed. Starting in 1984, the payroll tax rate for the self-employed rose sharply relative to that of employees. The 1983 reform also increased the FRA for individuals born in 1938 and later (see https://www.ssa.gov/planners/retire/retirechart.html). While benefits could still be claimed as early as age 62 with an actuarial reduction, the increase in the FRA effectively reduced the monthly benefit payable at each possible claiming age.[7] Thus, holding taxes and benefit levels at FRA constant, the increase in the FRA was equivalent to cutting the present value of lifetime benefits. Table 2 shows that only the CE data are appropriate for studying this reform.

Our third policy experiment comes from a 1981 reform that reduced dependents' benefits paid to children. Prior to August 1981, the children of deceased, retired, or disabled workers could receive Social Security dependent benefits if they were either under 18 or 18 to 21 and a full-time high school or college student. Legislation passed in August 1981 started the process of phasing out benefits for college students aged 18 to 21, with the phaseout complete by April 1985. Benefits for secondary school students older than 18 were eliminated by August 1982 (see DeWitt [2001] for additional details). This reform reduced the expected present value of the life insurance benefits available from Social Security for individuals with dependent children. Table 2 suggests that, again, only the CE data are appropriate for studying this reform.

1977 Reform: Models and Results

We use the PSID and NLS data to study the 1977 reform. The PSID is an ongoing panel survey of US households that began in 1968. The survey has been conducted each year until 1997 and every other year thereafter. We use data from 1968 to 1980. There are two family wealth variables available in 1968 to 1972 and again in 1975, 1979, and 1980.[8] First, families report whether they currently have savings equal to two months of income or more. Second, they report whether they have, at any point in the past five years, had savings equal to two months of income or more. The PSID also includes information on family money income and the ages of the head and spouse. We deflate income, expressing it in 1980 dollars, using the average Consumer Price Index for All Urban Consumers (CPI-U) over the survey year. We also divide income by its standard deviation, so the coefficients on income can be interpreted as the impact of a one standard deviation increase in income. In all our analysis, we use the PSID longitudinal family weights.

The NLS is a panel survey of men representative of the US population. We use the cohort of older men who were 45 to 59 in 1966, the first year of the surveys. Additional rounds of the survey continued either annually or biannually until 1983 and a final round was conducted in 1990. To measure savings, we use respondents' reported values of savings (including checking or savings accounts, or accounts with savings and loan companies or credit unions), bonds, and investments (which include bonds, stocks, and mutual funds). These values are reported in six different years, four years prior to the policy and two years after the policy. We also create a total savings variable, which is the sum of savings and investments. As with the PSID data, we deflate these variables using the CPI-U, expressing them in 1980 dollars. We also divide income by its standard deviation. In all the analyses, we use NLS custom longitudinal weights.

We begin by estimating the following difference-in-differences model:

$$y_{it} = \beta_1 \text{post}_t + \beta_2 \text{post}_t T_i + \beta_3 \text{income}_{it} + \beta_4 \text{age}_{it} + \gamma_t + \alpha_i + \varepsilon_{it}. \quad (10)$$

The dependent variable, y_{it}, is a measure of savings. In the PSID data, it is one of the indicators that family i has savings equal to two months or more of income in year t. In the NLS data, it is the values reported of checking/savings accounts, bonds, investments, and total savings. The independent variables are defined as follows: post_t is a dummy equal to 0 for years before 1977 and 1 for 1977 and later; T_i is a treatment dummy equal to 1 if the household head was born in 1917 or later and zero otherwise; income_{it} is a family money income; age_{it} is a set of dummies for the age of the head of household; γ_t is a year effect; α_i is an individual (NLS) or family (PSID) fixed effect; and ε_{it} is a stochastic error term. The coefficient of interest is β_2, which is expected to be positive. That is, our savings measures are expected to be higher for individuals whose benefits were cut following the 1977 reform. We would expect a higher income to cause higher savings. The age dummies control for the age profile of savings. The family fixed effects control for unobservable family heterogeneity, such as differences in the rate of time preference. The year dummies control for changes in macroeconomic conditions, differences in survey conditions, and other common factors that affect reported savings for all families in specific years.

Equation (10) does not take into account the degree of reduction in benefits that each post-1917 cohort experienced. Thus, we also estimate

$$y_{it} = \beta_1 \text{post}_t + \beta_2 \text{post}_t r_i + \beta_3 \text{income}_{it} + \beta_4 \text{age}_{it} + \gamma_t + \alpha_i + \varepsilon_{it}, \quad (11)$$

where r_i is the percentage reduction in monthly benefits payable at FRA experienced by individual i as a result of the 1977 amendments. Other variables are as defined above. The reduction in monthly benefits is 0 percent for individuals born before 1917, 13 percent for individuals born in 1917, 19 percent for individuals born in 1918, 26 percent for individuals born in 1919, and 29 percent for individuals born in 1920 or later.[9]

To investigate whether there is heterogeneity in responsiveness by income, we also estimate versions of equations (10) and (11) that include three-way interactions between income, the treatment dummy or reduction percentage, and the postperiod indicator. Our theoretical model predicts that the saving rates of higher-income households should be less responsive to benefit cuts due to the progressivity of the benefit formula. This theoretical argument applies directly to the PSID data, which measures savings relative to income. In the NLS regressions, the dependent variable is the dollar amount of savings. It is possible for the reforms to have a larger dollar impact on high-income households' savings even if the impact on the saving rate is smaller. In addition, for both data sets, it is possible that higher-income individuals may have a greater response as they are less likely to be liquidity constrained or eligible for means-tested welfare programs.

Table 3 shows the results from estimating equation (10) for the full PSID sample (column 1) as well as various subsamples. As identified in the last two rows of the table, the subsamples restrict the head of household's year of birth to 1910 to 1923 (the seven cohorts on either side of the reform), the survey year to 1972 to 1981, or both. The first four columns present results for the first measure of savings: whether the household currently has two months of income in savings. The next four columns present results for the second measure of savings: whether the household had two months of income saved at any point during the previous five years. None of the coefficients on the interactions between post-1977 and the treatment group indicator are statistically significant, though the point estimates are mostly positive as theory would predict. When three-way interactions including income are included, the signs of the post $\times$ treated interaction term again mostly go in the expected direction. For the first measure of savings, higher-income individuals appear to be less responsive to cuts in Social Security, consistent with the theoretical model. However, overall, only one of these coefficients is statistically significant.

Table 4 shows results for the PSID sample for equation (11), where the treatment group indicator is replaced with the percent reduction in benefits. Again, for the specifications where the post $\times$ treated dummy is interacted with income, the signs of the coefficients are mostly consistent with theory, particularly for the first savings measure. However, the coefficients are again mostly insignificant. Thus, we find at best weak evidence that the 1977 reform reduced the probability of having two months of income saved. Few coefficients are statistically significant, but the signs are generally consistent with the theoretical predictions and the size of the standard errors does not allow us to conclude that the reform had no effects.

Table 5 shows the results from estimating equation (10) using the NLS data. The interaction coefficient of interest is positive but small and insignificant for the savings regression. Moreover, the coefficients for the bonds, investments,

Table 3. Impact of Receiving Benefit Cut on Savings and Consumption (Panel Study of Income Dynamics).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	P (currently have two months' income saved)				*P* (had two months' income saved in past 5 years)			
	No interactions with income							
Post-1977 × treated	−0.00606	0.000113	0.0384	0.0603	0.00307	0.00787	0.0435	0.0623
	(0.0360)	(0.0617)	(0.0501)	(0.0684)	(0.0277)	(0.0489)	(0.0397)	(0.0554)
Family income	0.0647***	0.0674***	0.0540***	0.0836***	0.0244***	0.0234**	0.0215**	0.0321**
	(0.0132)	(0.0186)	(0.0177)	(0.0259)	(0.00716)	(0.0113)	(0.00999)	(0.0150)
	Interactions with income							
Post-1977 × treated	0.00604	0.0475	0.0408	0.125*	0.00582	0.00290	0.0370	0.0537
	(0.0390)	(0.0705)	(0.0527)	(0.0758)	(0.0305)	(0.0577)	(0.0422)	(0.0614)
Post-1977 × treated × family income	−0.00934	−0.0415	−0.00185	−0.0556**	−0.00212	0.00433	0.00512	0.00730
	(0.0130)	(0.0258)	(0.0137)	(0.0252)	(0.00882)	(0.0212)	(0.00997)	(0.0202)
Family income	0.0682***	0.0757***	0.0547***	0.0995***	0.0252***	0.0225*	0.0194*	0.0300*
	(0.0159)	(0.0192)	(0.0208)	(0.0260)	(0.00833)	(0.0123)	(0.0106)	(0.0169)
Observations	19,024	4,742	10,374	2,412	18,834	4,700	10,316	2,399
Number of families	2,601	604	2,601	604	2,601	604	2,601	604
Cohorts	All	1910–1923	All	1910–1923	All	1910–1923	All	1910–1923
Years	1968–1980	1968–1980	1972–1980	1972–1980	1968–1980	1968–1980	1972–1980	1972–1980

Note: All regressions also include year dummies, household head age dummies, and family fixed effects. Longitudinal weights are used. Standard errors clustered by family are in parentheses.
*p < .1.
**p < .05.
***p < .01.

Table 4. Impact of Percentage Reduction in Benefits on Savings and Consumption (Panel Study of Income Dynamics).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	P (currently have two months' income saved)				*P* (had two months' income saved in past 5 years)			
	No interactions with income							
Post-1977 × reduction	−0.0940	−0.179	0.0764	0.116	−0.0284	−0.104	0.137	0.186
	(0.128)	(0.248)	(0.187)	(0.301)	(0.102)	(0.197)	(0.153)	(0.253)
Family income	0.0647***	0.0672***	0.0540***	0.0837***	0.0244***	0.0232**	0.0216**	0.0324**
	(0.0133)	(0.0186)	(0.0177)	(0.0258)	(0.00717)	(0.0113)	(0.00999)	(0.0149)
	Interactions with income							
Post-1977 × reduction	−0.0494	0.0500	0.0827	0.405	−0.0146	−0.0997	0.113	0.151
	(0.141)	(0.299)	(0.199)	(0.349)	(0.113)	(0.243)	(0.163)	(0.287)
Family income × post-1977 × reduction	−0.0330	−0.181*	−0.00458	−0.219**	−0.0102	−0.00360	0.0170	0.0265
	(0.0457)	(0.0968)	(0.0482)	(0.0978)	(0.0308)	(0.0818)	(0.0352)	(0.0805)
Family income	0.0682***	0.0766***	0.0545***	0.1000***	0.0255***	0.0234*	0.0196*	0.0304*
	(0.0159)	(0.0191)	(0.0207)	(0.0258)	(0.00837)	(0.0121)	(0.0106)	(0.0168)
Observations	19,024	4,742	10,374	2,412	18,834	4,700	10,316	2,399
Number of families	2,601	604	2,601	604	2,601	604	2,601	604
Cohorts	All	1910–1923	All	1910–1923	All	1910–1923	All	1910–1923
Years	1968–1980	1968–1980	1972–1980	1972–1980	1968–1980	1968–1980	1972–1980	1972–1980

Note: All regressions also include year dummies, household head age dummies, and family fixed effects. Longitudinal weights are used. Standard errors clustered by family are in parentheses.
*p < .1.
**p < .05.
***p < .01.

Table 5. Impact of Receiving Benefit Cut on Savings (the National Longitudinal Survey).

	(1)	(2)	(3)	(4)
	Savings	Bonds	Investments	Total savings
	No interactions with income			
Post-1977 × treated	12.61	−92.02	−2,909	−3,722
	(1,921)	(348.9)	(4,266)	(4,913)
Income	6,018***	113.8	4,720***	11,822***
	(871.0)	(89.18)	(1,748)	(2,011)
	Interactions with income			
Post-1977 × treated	1,421	−360.0	−1,687	259.0
	(2,372)	(345.6)	(5,657)	(6,485)
Income × post-1977 × treated	−2,620	658.9**	−1,410	−2,946
	(2,131)	(296.0)	(7,673)	(8,096)
Income	3,065*	637.6***	5,997	13,343***
	(1,702)	(188.8)	(3,736)	(4,322)
Observations	18,740	20,577	20,621	21,006

Note: All regressions include year dummies, age dummies, and individual fixed effects. Survey years include 1966, 1969, 1971, 1976, 1981, and 1990. Robust standard errors are in parentheses and survey weights are used.
*$p < .1$.
**$p < .05$.
***$p < .01$.

and total savings regressions are all negative and insignificant. Three-way interactions with income produce mostly insignificant coefficients with mixed signs.

Table 6 shows the results from estimating equation (11) for the NLS data. The interaction coefficients for the savings specification is positive, but insignificant. Again, the coefficients on the bonds, investments, and total savings specifications are negative and insignificant. The three-way interactions with income are also largely similar for this specification. Overall, the evidence from the NLS suggests that people did not change their saving behavior in response to reduced Social Security benefits. However, the large standard errors do not allow us to conclude that the reform had no economically meaningful effects.

1983 Reform: Models and Results

The CE quarterly interview surveys are a rotating panel in which households report their income and expenditures in different categories for up to four

Table 6. Impact of Percentage Reduction in Benefits on Savings (the National Longitudinal Survey).

	(1)	(2)	(3)	(4)
	Savings	Bonds	Investments	Total savings
	No interactions with income			
Post-1977 × reduction	2,381	−74.42	−6,915	−9,140
	(7,333)	(1,389)	(17,006)	(19,285)
Income	6,018***	113.9	4,721***	11,822***
	(871.0)	(89.18)	(1,748)	(2,012)
	Interactions with income			
Post-1977 × reduction	6,893	−1,732	−2,016	8,005
	(9,180)	(1,374)	(23,839)	(26,533)
Income × post-1977 × reduction	−9,292	3,172***	−7,831	−16,334
	(8,523)	(1,045)	(34,630)	(36,002)
Income	3,037*	635.8***	5,713	12,774***
	(1,642)	(186.6)	(3,591)	(4,162)
Observations	18,740	20,577	20,621	21,006

Note: All regressions include year dummies, age dummies, and individual fixed effects. Survey years include 1966, 1969, 1971, 1976, 1981, and 1990. Robust standard errors are in parentheses and survey weights are used.
*$p < .1$.
**$p < .05$.
***$p < .01$.

consecutive quarters. We use the National Bureau of Economic Research (NBER) extracts of the CE. The NBER extracts aggregate the quarterly information provided by each household and member to create annual income and spending variables in a set of broad categories (see Harris and Sabelhaus 2000). Each survey includes a family-level file that includes aggregate consumption and income information for the family as well as a member-level file that includes demographic and earnings information for each member. We use data from 1980 through 1995.[10] We restrict the sample to families who are designated as "complete income reporters" and families who are interviewed for a full four quarters. We drop student households. We merge the family-level files to the individual-level files, and for each family, we retain the member with the highest earnings (defined as wages plus business and farm income) and define that individual as the household head.[11] We drop any families in which the household head has zero earnings, is not working, is working in the public sector, or is not the survey-designated head or spouse.[12]

As described by Harris and Sabelhaus (2000), we construct comprehensive measures of family consumption and before-tax family income by summing the various categories of consumption and income provided in the survey.[13] We also construct two measures of annual saving. The first defines saving as income minus taxes minus consumption. The second defines saving as the change in net worth.[14] We deflate these dollar amounts to 1980q1 dollars using the average CPI-U over the four quarters that each family is in the survey. We also divide income by its standard deviation.

We estimate the following difference-in-differences model:

$$s_{it} = \delta_1 \text{post}_t + \delta_2 \text{post}_t T_i + \delta_3 \text{income}_{it} + \delta_4 \text{age}_{it} + \mu_t + \varepsilon_{it}. \tag{12}$$

Here, s_{it} is defined as total saving for family i in year t, post_t is an indicator for years after 1983, T_i is a dummy for the treatment group (defined as either self-employed individuals, who experienced a tax increase, or individuals born in 1938 or later, who experienced an increase in FRA), age_{it} is a set of age dummies for the head of household, μ_t is a set of year dummies, and ε_{it} is a stochastic error term.

To determine the impact of the degree of tax increase or benefit reduction, we modify equation (12) as follows:

$$s_{it} = \delta_1 \text{taxrate}_{it} + \delta_2 \text{income}_{it} + \delta_3 \text{age}_{it} + \mu_t + \varepsilon_{it}, \tag{13}$$

$$s_{it} = \delta_1 \text{post}_t + \delta_2 \text{post}_t * \text{FRA}_{it} + \delta_3 \text{income}_{it} + \delta_4 \text{age}_{it} + \mu_t + \varepsilon_{it}. \tag{14}$$

Equation (13) replaces the first two terms on the right-hand-side of equation (12) with taxrate_{it}, the exact OASDI statutory tax rate (employer plus employee share) faced by the head of household. The tax rate varies in each year only according to the employment status (self-employed vs. wage earner) of the head of household. Equation (14) replaces the treated dummy in equation (12) with FRA_{it}, the exact postreform FRA faced by the head of household, thereby allowing larger increases in the FRA to have larger effects on saving. As we did for the 1977 reform, we also estimate versions of equations (12) to (14) that include interactions between the treatment $\times$ postmeasures and income.

For the payroll tax increase, we consider three samples: the full sample, families with heads aged 40 to 64 (i.e., prime working years), and families with heads born before 1938 (who were not affected by the FRA increase). For the FRA increase, we also consider three samples: the full sample, families with heads born between 1933 and 1942 (five years on each side of the 1938 cutoff), and families with heads who are private sector employees (all of whom experienced the same increase in the payroll tax). One shortcoming of the payroll tax experiment is that the relative increase in the payroll tax may have induced people to leave self-employment. Unfortunately, since the CE is not a panel survey, we cannot observe household heads' employment status before the reform.

Table 7 shows the impact of the payroll tax increase for the self-employed on saving (equation 12). In the first three columns, the dependent variable is total payroll (Social Security and Medicare) taxes paid by the household head as a share of the household head's earnings. Payroll taxes for wage earners are multiplied by two to account for the employer's contribution. The coefficient on the interaction term in these equations indicates the increase in the average payroll tax rate for the self-employed relative to wage earners following the 1983 reform. As figure 1 (which was based on statutory tax rates) suggests, observed tax rates indeed increased for the self-employed relative to wage earners.[15] Columns (4) to (6) suggest that following the 1983 reform, the self-employed reduced their saving by US$1,733 (if all age groups are included in the analysis) or US$2,381 (if only individuals aged 40 to 64 are included in the analysis) relative to wage earners. There is no statistically significant impact for individuals born before 1938, although the sample size is much smaller for this group. These three columns measure saving as disposable income minus consumption.[16] The other definition of saving suggests a more mixed picture—either no statistically significant change or a statistically significant increase (when only individuals aged 40 to 64 are included in the analysis). Thus, we have some evidence that the payroll tax increase reduced saving by one definition of saving. However, this finding is not robust across alternative definitions of saving. The coefficients on the three-way interactions with income mostly have positive signs, indicating that higher-income individuals may have been less responsive to the change.

Table 8 presents estimates of equation (13), in which the key independent variable is the exact statutory tax rate faced by the head of household. The results in columns (1) to (2) show that again, as expected, the observed average tax rate for the head of household is highly correlated with the statutory tax rate. The results in columns (3) to (9) are similar to those in table 6. They suggest that the reform reduced one measure of saving (disposable income minus consumption) for the self-employed, with mixed results for the other definition of saving. Between 1983 and 1984, the statutory payroll tax rate increased by 2.75 percentage points for the self-employed relative to wage earners. The results in column (3) suggest that this tax increase caused a US$1,676 (US$60,956 $\times$ 0.0275) decrease in saving for the full sample. The results in column (4) suggest that the tax

Table 7. Impact of Receiving Tax Increase on Saving (the Consumer Expenditure Surveys).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Variables		Tax rate			Income—taxes—consumption			Δ Assets	
					No interactions with income				
Self-employed	−0.0454***	−0.0448***	−0.0442***	−2,184***	−2,554**	−2,566**	729.4	−1,241	625.0
	(0.00118)	(0.00174)	(0.00160)	(704.9)	(1,082)	(1,026)	(894.5)	(1,492)	(1,433)
Self-employed × post-1983	0.0203***	0.0189***	0.0199***	−1,733**	−2,381*	114.2	800.5	4,032**	−2,156
	(0.00142)	(0.00215)	(0.00200)	(837.7)	(1,306)	(1,258)	(1,186)	(1,931)	(2,102)
Family income	−0.00960***	−0.0109***	−0.00897***	8,125***	7,946***	8,556***	2,554***	3,848***	5,487**
	(0.000448)	(0.000733)	(0.000384)	(214.1)	(326.9)	(396.9)	(896.2)	(1,483)	(2,527)
					Interactions with income				
Self-employed	−0.0493***	−0.0496***	−0.0491***	−3,058*	−1,196	−1,945	1,064	2,368	−61.72
	(0.00239)	(0.00370)	(0.00287)	(1,734)	(2,873)	(2,234)	(1,416)	(1,678)	(1,968)
Self-employed × post-1983	0.0292***	0.0295***	0.0271***	−2,158	−5,713*	−2,063	−1,186	−1,508	5,649
	(0.00296)	(0.00517)	(0.00346)	(1,921)	(3,142)	(2,645)	(2,528)	(4,086)	(6,211)
Self-employed × post-1983 × family income	−0.0103***	−0.0109***	−0.00966***	7,845***	7,610***	8,051***	1,253***	1,605***	1,963***
	(0.000715)	(0.000707)	(0.000749)	(299.0)	(390.4)	(414.4)	(378.1)	(599.2)	(714.8)
Family income	−0.00599***	−0.00649**	−0.00492**	141.6	2,001	1,350	1,135	3,358	−5,391
	(0.00170)	(0.00280)	(0.00205)	(1,649)	(2,569)	(2,286)	(2,029)	(2,982)	(4,919)
Observations	25,950	12,230	7,244	25,950	12,230	7,244	25,950	12,230	7,244
Sample restriction	All	Ages 40–64	Pre-1938	All	Ages 40–64	Pre-1938	All	Ages 40– 64	Pre-1938

Note: Dependent variable in columns (1) to (2) is household head's payroll taxes as a share of household head's earnings. Dependent variable in columns (3) to (6) is household saving in 1980q1 dollars. All regressions also include year dummies, household head age dummies, and household head cohort dummies. Attrition adjusted weights are used. Robust standard errors are in parentheses.

*$p < .1$.

**$p < .05$.

***$p < .01$.

Table 8. Impact of Statutory Tax Rate on Saving (the Consumer Expenditure Surveys).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Variables		Tax rate			Income—taxes—consumption			Δ Assets	
					No interactions with income				
Statutory tax rate	0.708***	0.653***	0.689***	−60,956**	−82,613*	2,613	31,293	141,108**	−72,966
	(0.0497)	(0.0754)	(0.0699)	(28,779)	(44,556)	(42,885)	(41,984)	(68,477)	(74,876)
Family income	−0.00961***	−0.0109***	−0.00897***	8,125***	7,946***	8,556***	2,554***	3,848***	5,487**
	(0.000447)	(0.000732)	(0.000384)	(214.1)	(326.9)	(396.9)	(896.3)	(1,483)	(2,527)
Self-employed	−0.0251***	−0.0260***	−0.0243***	−3,918***	−4,932***	−2,462***	1,547*	2,792**	−1,513
	(0.000809)	(0.00119)	(0.00121)	(448.0)	(717.3)	(699.6)	(804.5)	(1,280)	(1,759)
					Interactions with income				
Statutory tax rate	0.595***	0.485***	0.647***	−37,938	−66,790	−27,323	−51,857	−57,569	−361,454
	(0.0719)	(0.124)	(0.0899)	(35,077)	(53,770)	(54,989)	(119,008)	(213,755)	(338,579)
Family income	−0.0191***	−0.0240***	−0.0123***	10,066***	9,182**	6,211	−4,457	−11,670	−17,114
	(0.00396)	(0.00678)	(0.00424)	(2,742)	(4,215)	(4,374)	(7,975)	(13,378)	(21,298)
Statutory tax rate × family income	0.0805**	0.111**	0.0286	−16,516	−10,476	20,370	59,662	131,543	196,299
	(0.0313)	(0.0531)	(0.0373)	(23,813)	(36,614)	(38,910)	(74,779)	(124,910)	(205,555)
Self-employed	−0.0252***	−0.0261***	−0.0243***	−3,904***	−4,924***	−2,495***	1,497*	2,690**	−1,828
	(0.000815)	(0.00120)	(0.00121)	(448.3)	(718.0)	(704.1)	(832.1)	(1,330)	(1,983)
Observations	25,950	12,230	7,244	25,950	12,230	7,244	25,950	12,230	7,244
Sample restriction	All	Ages 40–64	Pre-1938	All	Ages 40–64	Pre-1938	All	Ages 40–64	Pre-1938

Note: Dependent variable in columns (1) to (2) is household head's payroll taxes as a share of household head's earnings. Dependent variable in columns (3) to (6) is household saving in 1980q1 dollars. All regressions also include year dummies, household head age dummies, and household head cohort dummies. Attrition adjusted weights are used. Robust standard errors are in parentheses.

*$p < .1$.

**$p < .05$.

***$p < .01$.

increase caused a US$2,272 (US$82,613 $\times$ 0.0275) decrease in saving for individuals aged 40 to 64. There is no evidence of an impact for individuals born before 1938. The three-way interactions with income are all insignificant.

Table 9 shows the results from estimating equation (12) for the FRA increase. The interactions between the indicator for being born in 1938 or later and the years after the reform have mixed signs and are statistically insignificant. When three-way interactions with income are included, the coefficients generally have signs that are consistent with theory but are mostly insignificant.

Table 10 shows similar results for equation (14), where the indicator for the treated group is replaced with the individual's FRA. Thus, we find little evidence that the reduction in lifetime benefits induced by the increase in the FRA led to an increase in saving. However, the standard errors are large, so we cannot rule out the possibility that there was an effect.

1981 Reform: Models and Results

We once again use the CE data to examine the impact of the elimination of student benefits by comparing life insurance premiums paid for households with and without dependent children. Using data on individual family members, we identify households with children who are living at home. Unfortunately, the data do not allow us to identify whether there are any children who are not living at home, such as college students who live in dorms. While it is college students who were most immediately affected by the elimination of student benefits, the reform still lowered the present value of the benefit paid to all children at or below college age. We estimate the following difference-in-differences model:

$$l_{it} = \delta_1 \text{transition}_t + \delta_2 \text{post}_t + \delta_3 \text{transition}_t \text{kids}_{it} + \delta_4 \text{post}_t \text{kids}_{it} + \delta_5 \text{income}_{it} + \delta_4 \text{age}_{it} + \mu_t + \varepsilon_{it}. \tag{15}$$

Here, l_{it} is life insurance premiums (in 1980q1 dollars) paid by household i in year t, transition_t is an indicator for time periods from the third quarter of 1981 through the first quarter of 1985 (the period over which student benefits were being phased out), post_t is an indicator for time periods after the first quarter of 1985 (when student benefits had been fully eliminated), and kids_{it} is an indicator for having kids under the age of 18 living at home.[17] If the student benefit elimination reduced life insurance holdings for affected families, we would expect δ_3 and δ_4 to be positive.

Table 9. Impact of Receiving FRA Increase on Saving (the Consumer Expenditure Surveys).

	(1)	(2)	(3)	(4)	(5)	(6)
Variables	Income—taxes—consumption			Δ Assets		
	No interactions with income					
Born in 1938 or later × post-1983	−45.68	778.8	−224.5	−1,101	302.8	−1,136
	(545.5)	(1,580)	(538.7)	(925.2)	(2,308)	(912.8)
Family income	8,024***	7,364***	7,914***	2,592***	2,190***	2,403**
	(213.8)	(377.7)	(233.5)	(889.1)	(781.1)	(1,101)
	Interactions with income					
Born in 1938 or later × post-1983	1,130	1,781	225.9	4,500	8,055**	7,797
	(1,139)	(2,725)	(1,149)	(4,817)	(3,529)	(6,515)
Family income	7,946***	5,387***	8,014***	2,124***	565.6	2,000***
	(520.0)	(1,628)	(420.1)	(651.6)	(1,278)	(697.7)
Born in 1938 or later × post-1983 × family income	−799.4	−610.3	−292.1	−3,542	−4,868**	−6,144
	(871.2)	(1,944)	(878.7)	(3,205)	(2,383)	(4,475)
Observations	25,950	4,662	23,287	25,950	4,662	23,287
Sample restriction	All	Born 1933–1942	Not self-employed	All	Born 1933–1942	Not self-employed

Note: Dependent variable is household saving in 1980q1 dollars. All regressions also include year dummies, household head age dummies, and household head cohort dummies. Attrition adjusted weights are used. Robust standard errors are in parentheses.
*p < .1.
**p < .05.
***p < .01.

Table 10. Impact of Full Retirement Age on Saving (the Consumer Expenditure Surveys).

	(1)	(2)	(3)	(4)	(5)	(6)
Variables	Income—taxes–consumption			Δ Assets		
			No interactions with income			
Full retirement age × post-1983	−2.904	228.0	−15.07	−92.04*	79.79	−70.04
	(31.81)	(299.4)	(30.93)	(51.67)	(432.7)	(49.00)
Family income	8,024***	7,367***	7,914***	2,590***	2,191***	2,402**
	(213.7)	(377.8)	(233.5)	(889.2)	(780.8)	(1,101)
			Interactions with income			
Full retirement age × post-1983	37.39	369.5	−5.945	149.7	973.5*	245.3
	(59.72)	(357.9)	(57.07)	(204.9)	(524.0)	(263.0)
family income	1,698	−97,975	19,389	79,471*	−121,683	69,807
	(39,803)	(158,403)	(35,890)	(41,362)	(145,966)	(43,919)
Full retirement age × post-1983 × family income	−32.63	−90.89	−4.545	−163.6	−580.3**	−238.2
	(56.42)	(229.4)	(53.49)	(152.6)	(284.9)	(199.5)
Observations	25,950	4,662	23,287	25,950	4,662	23,287
Sample restriction	All	Born 1933–1942	Not self-employed	All	Born 1933–1942	Not self-employed

Note: Dependent variable is household saving in 1980q1 dollars. All regressions also include year dummies, household head age dummies, and household head cohort dummies. Attrition adjusted weights are used. Robust standard errors are in parentheses.
*$p < .1$.
**$p < .05$.
***$p < .01$.

Because the reform primarily affected benefits for college students, life insurance value was primarily reduced for children who were likely to attend college. This allows us to add a third difference by using the head of household's education level to proxy for the likelihood of children under 18 attending college. We do this by estimating (6) with three-way interactions between time periods during and after the reform, having kids under 18, and an indicator for the head of household having some college education. Theory predicts that the coefficient on these three-way interactions will be positive.

Table 11 shows the results from estimating equation (15). In the first column, the coefficients on the interaction terms between having kids and being in the transition or postreform period have the expected positive sign. However, they are not individually or jointly significant. In the second column, the three-way interactions between having kids, being in the transition or postreform periods, and having some college education also have the expected positive sign. Furthermore, the coefficients on the two-way interactions between having kids and being in the transition or postreform periods are smaller in magnitude (and negative), and the three-way interactions with college have positive coefficients, consistent with the story that the reform had an effect for individuals with some college education but not those without. However, the coefficients on the three-way interaction terms are again not individually or jointly significant.[18] We conclude again that, while we cannot rule out the possibility that the reform affected life insurance purchases, we have no evidence to suggest that it did.

Discussion

As we have shown, the standard neoclassical model generates large predicted declines in savings from the introduction of Social Security in aggregate. Theory also predicts that the life insurance aspect of Social Security causes declines in life insurance holdings. On the other hand, our empirical analysis reveals that changes in Social Security policy do not have a statistically significant impact on savings or life insurance holdings. There are several reasons why such policy changes may generate less dramatic changes in observed savings than the predicted results from the model.

First, our standard errors are large. In our life insurance regressions, the coefficients of interest all have the expected signs; they are just not statistically significant. In all our regressions, our standard errors do not rule out the possibility of an economically meaningful effect that is consistent with theory. In recent decades, data sets like the HRS and SCF have started to

Table 11. Impact of Eliminating Student Benefits on Life Insurance Premiums Paid (the Consumer Expenditure Surveys).

	(1)	(2)
Variables	Life insurance premiums	
Have kids	87.31***	76.76**
	(24.65)	(31.11)
Have kids × (1981q3–1985q1)	20.35	−10.11
	(29.10)	(36.51)
Have kids × post-1985q1	9.047	−0.239
	(27.03)	(33.57)
College		27.36
		(43.38)
College × (1981q3–1985q1)		−59.38
		(46.35)
College × post-1985q1		−14.07
		(45.20)
College × have kids		34.42
		(55.78)
College × have kids × (1981q3–1985q1)		53.85
		(63.71)
College × have kids × post-1985q1		16.67
		(60.45)
Family income	0.0103***	0.0101***
	(0.000941)	(0.000983)
Observations	25,950	25,950

Note: Dependent variable is household life insurance premiums paid in 1980q1 dollars. All regressions also include year dummies, household head age dummies, and household head cohort dummies. Attrition adjusted weights are used. Robust standard errors are in parentheses.
*p < .1.
**p < .05.
***p < .01.

provide more detailed pictures of household wealth. In comparison, the data sets we use provide much less detail about saving behaviors or total savings, although they are the best available data sets that cover the period of the reforms we study. It is possible that the quality of the savings data from the 1970s and early 1980s is simply too noisy to precisely estimate these relationships.

Second, it is possible that individuals chose to delay retirement, rather than save more, in response to the benefit cuts studied. That may especially be the case for the 1983 reform, which cut lifetime benefits by raising the FRA. For example, Behaghel and Blau (2012) show that individuals tend to retire at whatever age is designated the FRA, possibly because they view this age as either a reference point or a recommendation by the government. On the other hand, Krueger and Pischke (1992) find little evidence that the 1977 reform had an impact on the labor force participation of affected older males.

Third, it is possible that individuals simply lack knowledge about their Social Security benefits and were unaware of the policy changes. While Smith and Couch (2014) show that younger workers are aware of the broad provisions of Social Security—for example, that it provides benefits to retirees and the families of deceased workers—Leibman and Luttmer (2015) find that many individuals are unaware of its specific design features. Gustman and Steinmeier (2004) further show that even older individuals make large errors in estimating their Social Security benefit levels. Thus, it is likely that many individuals are unaware of or do not pay attention to how changes in policy rules affect their benefits, particularly for small benefit changes. Lack of knowledge or attention is consistent with finding weak evidence of responsiveness to a tax increase but no evidence of responsiveness to a benefit cut. Since payroll taxes are deducted immediately from individuals' paychecks (or calculated and paid quarterly or annually for self-employed individuals), a tax increase is likely to be more salient and well-understood than a benefit cut.

Fourth, the timing of information about the policy change can be important in identifying any effects. If individuals learn about policy changes before they are implemented, then behavioral responses may not align with the timing of the policy change. As discussed by the Social Security Administration (n.d.), the problems created by the flawed indexation formula of the 1972 law became apparent almost immediately. The Social Security Trustees report of 1974 suggested that the program had a large long-term actuarial imbalance; by 1976, it had become "overwhelmingly clear" that reform would be required. Over this period, there was much public discussion about how to correct the flaw in the 1972 law. However, our identification strategy rests on the assumption that the 1977 reform affected different groups in different ways. And it is not clear to what extent the exact design of the reform—including which cohorts would be affected and by how much—could have been anticipated, as numerous alternative proposals were considered. Similarly, the financial shortfall that led to the 1983 reform was also apparent well in advance, and proposals to cut benefits were debated as early as 1981 (see Light 2005). Proposals to eliminate student benefits were considered as early as the late 1970s (DeWitt 2001). But here too, it is not clear to what extent the exact design of the reform could have been anticipated.

Fifth, individuals also likely lack information about whether the policy change is transitory or permanent. Households can easily observe that tax rates and benefit levels change frequently. For example, Shoven and Slavov (2006) document that promised internal rates of return vary considerably over each cohort's lifetime due to policy changes. Individuals may not respond much to any one change if they are not sure how long that change will last.

Sixth, expectations may also affect the timing of behavioral responses. Expectations could matter in several ways. First, any policy change should be taken in context of the overall Social Security budget. Our theoretical model assumes a balanced government budget. However, our empirical work relies on benefit cuts and tax increases designed to restore fiscal balance. In an environment where Ricardian equivalence holds, the timing of policy changes to restore long-run fiscal balance to the Social Security system should not influence behavior if it does not change individual expectations about future policy. As discussed above, the actuarial shortfalls that triggered the 1977 and 1983 reforms were apparent well in advance, although it is not clear to what extent the differences in impact across individuals (on which our identification rests) could have been anticipated. Alternately, Dominitz, Manski, and Heinz (2003) argue that many individuals do not expect to receive any Social Security benefits, so changes in rules may not influence behavior. More generally, several articles including Gomes, Kotlikoff, and Viceira (2007), van der Wiel (2008), and Caliendo, Gorry, and Slavov (2016) study the impact of uncertainty on saving decisions and welfare.

Finally, a large fraction of the population may not save at all due to liquidity constraints or means-tested welfare programs (see Hubbard, Skinner, and Zeldes 1995). These individuals are less likely to be responsive to changes in Social Security wealth. To the extent that liquidity constraints and utilization of means-tested programs are correlated with income, this story predicts that higher-income families should be more responsive to the policy changes. In our empirical models, however, we found very little statistically significant evidence of heterogeneity in responsiveness by income, and point estimates often suggested that higher-income individuals were less responsive.

Conclusions

In this article, we review the theoretical implications of Social Security on private savings, establishing that Social Security strongly crowds out private saving behavior. With these theoretical predictions in mind, we empirically evaluate the effect of Social Security on savings and insurance purchasing behavior using three different policy changes. Despite the strong theoretical predictions, we find little evidence to support that Social Security crowds out private savings. We posit that lack of knowledge about Social Security program details in general or the specific implications of the policy change could mitigate any effect of the policy on individual saving behavior.

Acknowledgments

We would like to thank Andrew Biggs and participants in the Consumer Expenditure Survey Microdata Users' Workshop and the National Tax Association Annual Conference on Taxation for their helpful comments. We also thank Muzhdah Karimi for excellent research assistance.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by the Mercatus Center at George Mason University. The findings and conclusions expressed are solely those of the authors.

Notes

1. We do not model the individual's labor supply decisions for two reasons. First, correctly modeling the labor supply choice requires a fully calibrated model that includes the age-based productivity profile by wage type, the Social Security–delayed retirement credit (and how this dimension of Social Security has evolved significantly over time), opportunities for part time work during retirement as well as stochastic factors such as health, disability, and employment shocks. Second, we do not run any empirical tests on the effect of Social Security on labor supply decisions. While not including labor supply allows us to focus on saving decisions, doing so is a limitation because we miss how labor supply and saving decisions potentially interact.

2. In so far as bequest income is concerned, some macroeconomic models assume that the assets of the deceased are spread evenly across all cohorts of the living. Doing this in our model would be a minor adjustment. Beyond this, there are two major modeling adjustments that would make the model significantly more complex but might generate new insights about the impact of Social Security on savings. First, while we assume that wage income is independent and identially distributed across generations, one could instead assume that it is persistent and, therefore, a high-wage parent bequeaths her or his wealth to a high-wage child. Second, one could build specific intergenerational linkages into the model, so that bequest income received is a function of the timing of the deaths of an individual's entire ancestry as in Caliendo, Guo, and Hosseini (2014).
3. This survival function is estimated from Social Security Administration cohort mortality tables for men born in 1990.
4. Recent changes to the full retirement age (FRA) and the delayed retirement credit were set in motion by the 1983 reform and therefore fully anticipated. The elimination of the earnings test after FRA in 2000 did not cause a major change in the lifetime value of benefits, as any benefits withheld due to the earnings test are paid—with an actuarial adjustment—once the individual reaches the age at which the earnings test no longer applies.
5. It is possible for Social Security wealth to have an impact on portfolio composition in addition to the total level of savings. Since Social Security wealth is a relatively safe asset, a reduction in it may induce individuals to reallocate their portfolios toward other safe assets. See Delavande and Rohwedder (2011) and Yogo (2016) for further discussion. However, we focus on the total level of savings both because that is in line with the theoretical model and because the available data sets generally do not have good measures of portfolio composition.
6. The error in the indexation formula effectively double-indexed benefits for individuals who had not yet claimed in 1972. Thus, double indexation primarily affected individuals born between 1910 and 1916, who turned 62 in 1972 or later but were not affected by the 1977 reform. We do not use the 1972 error as a policy experiment as individuals born before 1910 received ad hoc cost-of-living adjustments. See Social Security Administration (2004) and Krueger and Pischke (1992) for additional details.
7. The 1983 reform also gradually increased the delayed retirement credit, or the actuarial increase in the monthly benefit for each year of delayed claiming beyond FRA. The increases in the delayed retirement credit were phased in for individuals born in 1925 and later. However, very few individuals delay benefits beyond FRA (see, e.g., Goda et al. 2017). Also, the combination of the increases in the delayed retirement credit and the FRA in most cases reduces the

benefit that individuals born in 1938 and later can receive regardless of the age at which they claim. For additional details, see https://www.ssa.gov/oact/ProgData/ar_drc.html.

8. These wealth variables are not available beyond 1980, although more detailed wealth variables are consistently available starting in the 1990s.
9. We obtain the percentage reduction in benefits from https://www.ssa.gov/history/notchfile3.html.
10. Data are not available on families entering the survey during the third and fourth quarters of 1985 and 1995.
11. If the high earner is not unique, we order the high earners according to their relationship to the survey-designated household head as follows: head, spouse, child, grandchild, in-law, brother/sister, mother/father, other relative, and unrelated individual. We then retain the first high earner. In more than 80 percent of families, the high earner is the survey-designated household head. In almost 95 percent of families, the high earner is the survey-designated household head or spouse.
12. Under the 1983 reform, federal employees began to be covered by Social Security. However, we cannot distinguish federal employees from state or local employees. Many state and local employees are covered by Social Security, but others are not.
13. For income, we add up wages, business income, farm income, rent, dividends, interest, pension income, Social Security benefits, Supplementary Security Income, unemployment compensation, workers compensation, welfare payments, scholarships, food stamps, contributions from others (including alimony and child support), lump sums (e.g., from inheritances), and insurance refunds. We then subtract contributions made to others (including alimony and child support).
14. The change in net worth sums pension and retirement contributions, changes in checking and savings account balances, changes in stocks and bonds, investments made in a farm or business, net properties purchased, additions and alterations made to properties, and the net reduction in debt.
15. Observed tax rates differ from statutory tax rates because they also include contributions to Medicare and because the measure of earnings reported in the data may not correspond exactly to income that was subject to the payroll tax, possibly due to the payroll tax cap, noncompliance, or measurement error.
16. Since taxes—including payroll taxes—are subtracted from gross income to arrive at disposable income, and since payroll tax rates increased for the self-employed following the reform, one might be concerned that there is a mechanical relationship between the savings definition and the reform. But this

specification is consistent with the theory, which predicts that the payroll tax increase reduces saving rather than consumption.

17. Having children aged 18 to 21 living at home is relatively uncommon. Results are largely similar if $kids_{it}$ is defined as having kids aged 21 and younger living at home.
18. Three-way interactions with income also have insignificant coefficients. Results are available upon request.

References

Aguila, Emma. 2011. "Personal Retirement Accounts and Saving." *American Economic Journal: Economic Policy* 3(4): 1–24.

Alonso-Ortiz, Jorge. 2014. "Social Security and Retirement across the OECD." *Journal of Economic Dynamics and Control* 47:300–16.

Attanasio, Orazio P., and Agar Brugiavini. 2003. "Social Security and Households' Saving." *Quarterly Journal of Economics* 118(3): 1075–119.

Attanasio, Orazio P., and Susann Rohwedder. 2003. "Pension Wealth and Household Saving: Evidence from Pension Reforms in the United Kingdom." *American Economic Review* 93(5): 1499–521.

Bagchi, Shantanu. 2016. "Is the Social Security Crisis Really as Bad as We Think?" *Macroeconomic Dynamics* 20(3): 737–76.

Behaghel, Luc, and David M. Blau. 2012. "Framing Social Security Reform: Behavioral Responses to Changes in the Full Retirement Age." *American Economic Journal: Economic Policy* 4(4): 41–67.

Bernheim, B. Douglas. 1987. "The Economic Effects of Social Security: Toward a Reconciliation of Theory and Measurement." *Journal of Public Economics* 33(3): 273–304.

Bernheim, B. Douglas, and Lawrence Levin. 1989. "Social Security and Personal Saving: An Analysis of Expectations." *American Economic Review* 79(2): 97–102.

Blau, David M. 2016. "Pensions, Household Saving, and Welfare: A Dynamic Analysis of Crowd Out." *Quantitative Economics* 7(1): 193–224.

Bütler, Monika. 2001. "Neoclassical Life-cycle Consumption: A Textbook Example." *Economic Theory* 17:209–21.

Caliendo, Frank N., Aspen Gorry, and Sita Slavov. 2016. "The Cost of Uncertainty about the Timing of Social Security Reform." NBER Working Paper No. 21585. Cambridge, MA: National Bureau of Economic Research.

Caliendo, Frank N., Nick Guo, and Roozbeh Hosseini. 2014. "Social Security Is NOT a Substitute for Annuity Markets." *Review of Economic Dynamics* 17: 739–55.

Conesa, Juan C., and Carlos Garriga. 2008. “Optimal Fiscal Policy in the Design of Social Security Reforms.” *International Economic Review* 49(1): 291–318.

Delavande, Adeline, and Susann Rohwedder. 2011. “Individuals’ Uncertainty about Future Social Security Benefits and Portfolio Choice.” *Journal of Applied Econometrics* 26(3): 498–519.

DeWitt, Larry. 2001. “The History of Social Security ‘Student” Benefits.’ Research Note #11, Social Security Administration Historian’s Office. Washington, DC: Social Security Administration.

Diamond, Peter A., and Jerry A. Hausman. 1984. “Individual Retirement and Savings Behavior.” *Journal of Public Economics* 23(1–2): 81–114.

Dominitz, Jeff, Charles F. Manski, and Jordan Heinz. 2003. “Will Social Security Be There for You?: How Americans Perceive their Benefits.” NBER Working Paper No. 9798. Cambridge, MA: National Bureau of Economic Research.

Engelhardt, Gary V., and Anil Kumar. 2011. “Pensions and Household Wealth Accumulation.” *Journal of Human Resources* 46(1): 203–36.

Feldstein, Martin. 1974. “Social Security, Induced Retirement, and Aggregate Capital Accumulation.” *Journal of Political Economy* 82(5): 905–26.

Feldstein, Martin. 1996. “Social Security and Saving: New Time Series Evidence.” *National Tax Journal* 49(2): 151–64.

Feldstein, Martin, and Anthony Pellechio. 1979. “Social Security and Household Wealth Accumulation: New Microeconometric Evidence.” *Review of Economics and Statistics* 61(3): 361–68.

Feng, Jin, Lixin He, and Hiroshi Sato. 2011. “Public Pension and Household Saving: Evidence from Urban China.” *Journal of Comparative Economics* 39(4): 470–85.

Franco, Annie, Neil Malhotra, and Gabor Simonovits. 2014. “Publication Bias in the Social Sciences: Unlocking the File Drawer.” *Science* 345(6203): 1502–05.

Gale, William G. 1998. “The Effects of Pensions on Household Wealth: A Reevaluation of Theory and Evidence.” *Journal of Political Economy* 106(4): 706–23.

Goda, Gopi Shah, Shanthi Ramnath, John B. Shoven, and Sita Nataraj Slavov. 2017. “The Financial Feasibility of Delaying Social Security: Evidence from Administrative Tax Data.” *Journal of Pension Economics and Finance* (forthcoming).

Gokhale, Jagadeesh, Laurence J. Kotlikoff, and John Sabelhaus. 1996. “Understanding the Postwar Decline in U.S. Saving: A Cohort Analysis.” *Brookings Papers on Economic Activity* 1:315–407.

Gomes, Francisco J., Laurence J. Kotlikoff, and Luis M. Viceira. 2007. “The Excess Burden of Government Indecision.” NBER Working Paper No. 12859. Cambridge, MA: National Bureau of Economic Research.

Gustman, Alan L., and Thomas L. Steinmeier. 1998. “Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study.” NBER Working Paper No. 6681. Cambridge, MA: National Bureau of Economic Research.

Gustman, Alan L., and Thomas L. Steinmeier. 2004. "What People Don't Know about Their Pensions and Social Security." In *Private Pensions and Public Policies*, edited by William G. Gale, John B. Shoven, and Mark J. Warshawsky, 57–125. Washington, DC: Brookings Institution.

Harris, Ed, and John Sabelhaus. 2000. *Consumer Expenditure Survey. Family Level Extracts 1980:1-1998:2*. Cambridge, MA: National Bureau of Economic Research, Unpublished Manuscript.

Hong, Jay H., and José-Víctor Ríos-Rull. 2007. "Social Security, Life Insurance and Annuities for Families." *Journal of Monetary Economics* 54(1): 118–40.

Hosseini, Roozbeh. 2015. "Adverse Selection in the Annuity Market and the Role for Social Security." *Journal of Political Economy* 123:941–84.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103:360–99.

Kitao, Sagiri. 2014. "Sustainable Social Security: Four Options." *Review of Economic Dynamics* 17(4): 756–79.

Kotlikoff, Laurence J. 1979. "Testing the Theory of Social Security and Life Cycle Accumulation." *American Economic Review* 69(3): 396–410.

Krueger, Alan B., and Jörn-Steffen Pischke. 1992. "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation." *Journal of Labor Economics* 10(4): 412–37.

Leibman, Jeffrey B., and Erzo F. P. Luttmer. 2015. "Would People Behave Differently if They Better Understood Social Security? Evidence from a Field Experiment." *American Economic Journal: Economic Policy* 7(1): 275–99.

Leimer, Dean R., and Selig D. Lesnoy. 1982. "Social Security and Private Saving: New Time Series Evidence." *Journal of Political Economy* 90(3): 606–29.

Light, Paul C. 2005. "The Crisis Last Time: Social Security Reform." Brookings Institution, March 5. Accessed May 18, 2017. https://www.brookings.edu/opinions/the-crisis-last-time-social-security-reform/.

Li, Yue. 2016. "Social Security Survivors Benefits and Life Insurance." Working paper, Albany, NY: University of Albany.

Parker, Jonathan. 1999. "Spendthrift in America? On Two Decades of Decline in the U.S. Saving Rate." In *NBER Macroeconomics Annual*, edited by Ben Bernanke and Julio Rotemberg, vol. 14, 317–70. Cambridge, MA: MIT Press.

Shoven, John B., and Sita Nataraj Slavov. 2006. "Political Risk Versus Market Risk in Social Security." NBER Working Paper No. 12135. Cambridge, MA: National Bureau of Economic Research.

Smith, Barbara A., and Kenneth C. Couch. 2014. "How Effective Is the Social Security Statement? Informing Younger Workers about Social Security." *Social Security Bulletin* 74(4): 1–19.

Social Security Administration. 2004. “The ‘Notch’ Provision.” SSA Publication No. 05-10042. Washington, DC: Social Security Administration.
Social Security Administration. (n.d.), “The Commission on the Social Security ‘Notch’ Issue.” Accessed May 18, 2017. https://www.ssa.gov/history/notchfile1.html.
Van der Wiel, Karen. 2008. “Preparing for Policy Changes: Social Security Expectations and Pension Scheme Participation.” Institute for the Study of Labor, IZA DP No. 3623. Bonn, Germany: Institute for the Study of Labor.
Yaari, Menachem. 1965. “Uncertain Lifetime, Life Insurance and the Theory of the Consumer.” *Review of Economic Studies* 32(2): 137–50.
Yogo, Motohiro. 2016. “Portfolio Choice in Retirement: Health Risk and the Demand for Annuities, Housing, and Risky Assets.” *Journal of Monetary Economics* 80(Supplement C): 17–34.

Author Biographies

Sita Slavov is a professor of public policy at George Mason University’s Schar School of Policy and Government, and a faculty research fellow at the National Bureau of Economic Research. She received her PhD in economics from Stanford University and has previously worked at the American Enterprise Institute, the Council of Economic Advisers, and Occidental College.

Devon Gorry is an assistant professor in the John E. Walker Department of economics at Clemson University. She is an applied microeconomist who studies issues related to health, education, and labor. She received her PhD in economics from the University of Chicago.

Aspen Gorry is an assistant professor in the John E. Walker Department of Economics at Clemson University. He received his PhD in economics from the University of Chicago and has previously worked at the University of California, Santa Cruz, American Enterprise Institute, and Utah State University.

Frank N. Caliendo is the department head of economics and finance in the Jon M. Huntsman School of Business at Utah State University. He is also the academic director of the Center for Growth and Opportunity at Utah State University.