

SOCIAL SECURITY AND UNSECURED DEBT *

Erik Hurst[†]

Graduate School of Business
University of Chicago
and NBER

Paul Willen[‡]

Federal Reserve Bank of Boston
and NBER

First version: July 2003

This version: June 2006

*We thank Charlie Brown, Jeff Liebman, Borja Larrain, Olivia Mitchell, Geoff Tootell, Kent Smetters, Steve Zeldes, two anonymous referees, and seminar participants at the Cleveland Fed, the Chicago GSB Macro Lunch, Northwestern, U. Penn, and the NBER Summer Institute Social Security Working Group for helpful comments and suggestions and Francisco Torralba for excellent research assistance and helpful comments and suggestions. Both authors gratefully acknowledge research support from the Graduate School of Business at the University of Chicago.

[†]Phone: (773) 834-4073. E-mail: erik.hurst@gsb.uchicago.edu.

[‡]Phone: (617) 973-3149. E-mail: paul.willen@bos.frb.org

Abstract

In this paper, we use a calibrated life-cycle model to explore quantitatively ways of reducing the burden of social security in a world populated by both optimizing and rule-of-thumb consumers. Social security contributions force young households with upward-sloping income profiles to save a sizeable portion of their income for retirement, when their optimal consumption plan would likely have them either saving little or borrowing. We first use household data to document that young households have accumulated social security contributions that are large relative to debt holdings. Then, using a calibrated life-cycle model, we show that both allowing households to use social security wealth to pay off their debt and exempting young households from social security contributions (but in both cases requiring higher contributions later) mitigate many of the inefficiencies of social security from the perspective of life-cycle financial planning. Specifically, in our preferred experiment, which exempts households whose heads are under 30 from making social security contributions, we find that certainty-equivalent consumption increases by 3.4 percent for optimizing households and by 3.3 percent for rule-of-thumb households.

1 Introduction

Many households, particularly the young, hold large amounts of unsecured debt on which they pay a high interest rate.¹ At the same time, many of these households have made substantial contributions into the social security system on which they earn a low return, thus imposing a potentially suboptimal financial planning regime. In this paper, we use a calibrated life-cycle model to explore quantitatively ways of reducing the burden of social security on such households. We show that both allowing households to use social security wealth to pay off their debt and exempting young households from making social security contributions (but in both cases requiring higher contributions later so as not to affect the present value of total lifetime contributions) can mitigate many of the inefficiencies of social security from the perspective of life-cycle financial planning.

We begin the paper by comparing empirically the distribution across individuals of non-collateralized debt and accumulated social security contributions. In our sample of households under the age of 40, 62 percent have unsecured debt. We show that if households could access their accumulated social security contributions to pay off debt, only 17 percent would continue to hold debt. And for that 17 percent, total debt would be dramatically reduced: for the 90th-percentile household in the debt distribution, unsecured debt would fall from 84 percent to 33 percent of that household's average income. We conclude that there is a large potential margin to reduce household debt that could be achieved if young households were allowed access to their past social security contributions.

In Section 3, we construct a dynamic life-cycle portfolio choice model with an embedded social security program. We follow Campbell and Mankiw (1989) and others in assuming that the world is populated by two types of households: optimizing households that use financial assets to maximize utility and "rule-of-thumb" households that simply set consumption equal to income. Evidence from the consumption literature suggests that households break roughly into these two categories, especially with respect to retirement planning. Rule-of-thumb households also provide a justification for the existence of social security. Specifically, social security prevents destitution for rule-of-thumb households once they enter retirement and have no income and no savings. For the optimizing households, we adapt the model developed by Davis, Kubler, and Willen (2006). In that model, households can invest in stocks and bonds and can also take out unsecured loans. We specify that the interest rate on unsecured debt (that is, the borrowing rate) exceeds the interest rate on bonds (that

¹In general, throughout this paper, the age of a household is defined by the age of the household head. By young here, we mean, households with a head whose age is less than 40.

is, the lending rate). Such an assumption is consistent with the empirical pattern of observed borrowing and lending rates.² The parameterization of this model roughly matches the life-cycle borrowing behavior documented in Section 2.

In Section 4, we analyze the effects of two policy experiments aimed at alleviating the inefficiency to young households that has them simultaneously making large social security contributions and holding unsecured debt on which they pay a high interest rate. In our first experiment, we allow households currently in the social security system to access their accumulated social security contributions to pay off debt. In our second experiment, we build on an idea of Hubbard and Judd (1987) and exempt young households from social security contributions. Under both proposals, households would contribute more to social security (via higher taxes) later in their working lives to compensate for their reduced contributions while young. Such an assumption ensures unchanged social security benefits upon retirement. Both of the above proposals lead to increases in saving, reductions in debt, and increases in certainty-equivalent consumption.³ For example, under our preferred policy, we move households from our existing system to one in which households under 30 are exempt from contributing to social security. We find that such a policy raises certainty-equivalent consumption by 3.4 percent for optimizing households and 3.3 percent for rule-of-thumb households.

Our analysis also yields valuable insights for the current debate on reforming social security. One often-discussed inefficiency of the social security system is that it forces everyone to invest all of their social security contributions in low-yielding assets. We show that our proposed social security exemptions raise welfare by as much as, or, in some cases, more than the result that would be achieved if the social security system could achieve a 5 percent real riskless return on its investments. Thus, exemptions allow us to address one additional criticism of the current social security system without the costly and politically problematic option of allowing individuals to invest social security funds in risky assets.

Finally, we draw attention to one other possibility offered by our experiments. By either giving an exemption or allowing a withdrawal but requiring higher contributions later, the government is effectively lending money to households. We call the interest rate on such loans the internal borrowing rate (IBR). The conditions we impose imply that the IBR equals the internal rate of return on social security investment. But the welfare gains from these “loans” are so big that the government could charge a higher IBR and still make households better off. We show that a combination of the

²See Davis et al., 2006.

³ Certainty-equivalent consumption is the level of constant consumption over the life cycle that yields the same utility as the risky stream the household actually receives.

age-40 exemption and an increase in the IBR from 2 percent to 5 percent still leads to a significant increase in welfare for both optimizing and rule-of-thumb households. We also show that raising the IBR would significantly improve the finances of the social security system. In other words, we show that the government could borrow at 2 percent, lend at 5 percent, and still make households better off. Therefore, one version of our proposal could improve the solvency of the social security system while at the same time making households better off.

This paper adds to the literature that examines the effects of a social security program on household behavior. For a survey, see Feldstein and Liebman (2002). We draw the reader's attention to three particularly relevant papers. Feldstein and Rangelova (2001) calculate the effects on retirement income of their proposal to replace the current pay-as-you-go social security system with individual accounts in which households can invest in equity. They find that people typically do better with individual accounts, although there is a small probability that they will do worse.

Our approach to reforming social security owes a significant debt to Hubbard and Judd's (1987) analysis of social security, in which households face borrowing constraints. They show, as we do, that an exemption from social security contributions early in the life cycle yields substantial welfare benefits. Our analysis differs from theirs in two key respects. First, we analyze the effects of exemptions on households that do not plan for the future, our "rule-of-thumb" households. And second, households in our model can invest in stocks and borrow, activities not offered to households in Hubbard and Judd's model.

Campbell, Cocco, Gomes, and Maenhout (2001) explore the effects on portfolio choice and utility of allowing households to invest some of their social security wealth in equities, in the context of a life-cycle model with no borrowing. They find that a combination of smaller contributions and investment in equities leads to substantial welfare increases.

1.1 Some caveats

Before continuing with the paper, we draw the reader's attention to several important aspects of our analysis.

First, we make no assumptions about how social security is funded. Nor do we take any stand on changing the financing of the social security system. Our main vantage point is that of the individual household, which contributes money to social security while working and receives benefits in retirement with certainty. We use the internal rate of return estimated by Leimer (1994) to link contributions with benefits, but we do not assume that the social security system invests the contributions in an

individual account or in a pension fund. However, we measure the effects of various social security schemes on the difference between annual contributions by workers and annual retiree benefits at various points in time — in other words, we evaluate how our proposals would affect the solvency of a pay-as-you-go system. It should be stressed again that our policy experiments are designed to leave the benefit portion of social security unchanged.

Second, we assume that households do not face longevity risk. This is an important omission, as some researchers argue that by providing an indexed life annuity, social security allows households to manage longevity risk.⁴ Further, they argue that markets fail to provide such annuities. This benefit of social security is missing from our model. However, none of our proposed schemes — except our “straw man” of eliminating social security altogether — has any effect on the retirement portion of social security. Thus, the benefits shown in this paper are incremental to the social benefits of a mandated, indexed life annuity.

Third, in all of our policy experiments, we require that households contribute at least as much in present-value terms and receive exactly the same benefits as they do in the current system. For example, we consider allowing households to withdraw their wealth from social security to pay off debt. But we do not allow them to “opt out” of social security at all — in fact, their subsequent contributions go up. So our experiment is completely different from, for example, that of Smetters and Walliser (2004), who allow households to leave social security entirely. In addition, our model preserves the “commitment” aspect of social security. Some researchers (Akerlof 1998, for example) have argued that households do not trust themselves to save, and that therefore they vote for a government program that compels them to do so. The options we consider change only the life-cycle structure of social security, not the level of contributions. Social security will still require that households save enough to guarantee an income in retirement equal to 43 percent of their average working income.

Fourth, the main point of the paper — that the ideal life-cycle profile of contributions is not flat — applies equally well to any tax. Given the choice, liquidity-constrained households with a hump-shaped income profile would rather pay less income tax when young and more income tax when middle-aged. We focus on social security for two reasons. First the explicit purpose of social security, unlike that of income taxes, is to smooth life cycle consumption. Second, a progressive income tax approximates the ideal life-cycle structure by lowering tax rates when income is

⁴See Abel (1986) and Eckstein, Eichenbaum, and Peled (1985), among others. In our model, the rationale for the social security system is to prevent destitution for rule-of-thumb households who saved too little to sustain their consumption in retirement.

low. Since social security taxation is, in fact, regressive, it is a natural target for our analysis.

Finally, we draw the reader to two limitations of our analysis. First, we assume that the labor supply evolves independently of taxes on labor income, an assumption that presents particular problems for those of our policy experiments that introduce major changes in both the level of social security taxes and their life-cycle profile. For example, compared with the existing system, the age-50 exemption leads to a significant reduction in taxes before age 50 and a huge increase in taxes after age 50. Such a change in the tax code could lead to major changes in life-cycle labor supply.⁵ However, our preferred model of exempting households up to the age of 30 (as opposed to the age-40 or age-50 exemptions) results in large welfare gains and requires only a small increase in the marginal tax rate of the average household after the age of 30 — from 10.6 percent to 12.9 percent. Relative to changes in the social security tax rate observed over the last quarter century, this change is very small and would have minimal effects on household labor supply.

The second limitation of our analysis is that we have constructed a partial-equilibrium model. Obviously, if a policy change has significant partial-equilibrium effects, one would imagine that it would have significant general-equilibrium effects. Researchers have found, for example, that investing some of the social security trust fund in equities — which we modeled as an increase in the internal lending rate — would have significant macroeconomic effects.⁶

2 Empirical facts about borrowing and social security

The PSID is a large, nationally representative survey, started in 1968, that tracks social and economic variables of a given household over time. Each year, the survey collects demographic information such as age, race, family composition, and education levels of all members in each household. Among other information, individuals report their labor market participation and earned labor income.

On occasion, the PSID supplements the main data set with special modules. In 1984, 1989, 1994, and 1999, the PSID asked households extensive questions about their wealth. Specifically, households report holdings of cash, stocks, bonds, mutual

⁵Given that the social security benefits paid during retirement are tied to earned income during working years, the distortion caused to labor supply from raising social security taxes on older working households will be mitigated.

⁶See Bohn (1999) and Diamond and Geanakoplos (2003) for examples.

funds, saving accounts, checking accounts, government savings bonds, Treasury bills, Individual Retirement Accounts, bond funds, cash value of life insurance policies, valuable collections for investment purposes, and rights in a trust or estate. Additionally, respondents report the value of their main home, the value of their mortgage debt outstanding, and their net positions in other real estate, businesses, and vehicle ownership. Of particular interest to this study are the respondents' reports of their holdings of unsecured debt (including store and credit card debt, student loans, and other personal loans).

Given that the PSID measures wealth (and indebtedness) at five-year intervals starting in 1984, we focus our analysis on the 1999 wealth supplement. There are two reasons for this. First, in order to compute social security wealth, we need a long history of earnings for each individual in the survey. Using the 1999 data thus allows us as many as 31 annual observations of individual earnings. Second, innovation in financial markets resulted in an explosion in credit card use between the late 1980s and the late 1990s. Focusing on more recent time periods, therefore, may provide a more representative picture of a household's steady-state holdings of unsecured debt. However, for completeness, we repeated our analysis for a sample of 1989 households, and our main conclusions were unchanged.

Our main sample included all household heads in the 1999 PSID between the ages of 22 and 39. We focused on younger households since most unsecured borrowing occurs among this group. Unfortunately, we do not have full earnings histories for all household heads in the sample. To see why, note that the PSID tracks only core PSID members over time. A core PSID member is either an original respondent from 1968 (when the survey started) or a descendent of an original sample member. Additionally, given PSID definitions, the male is always classified as the household head for all married and cohabiting couples. This combination results in some male heads having earnings histories extending back only to the date they married into the survey. Rather than imputing the missing income histories for these household heads, we restricted our sample to include 1) all single-headed households (both male and female) and 2) all married households where the head has a complete earnings history. Given that all the households in the 1999 PSID under the age of 40 are descendants who are equally likely to be men or women, our restriction does not bias our analysis in any way. The resulting sample size is 2,077 households.⁷

While the PSID does not directly measure social security contributions, we can compute a measure of total social security contributions using households' detailed

⁷See Appendix Table A1 for descriptive statistics of the sample.

earnings histories coupled with the social security tax tables.⁸ Specifically, we calculate the social security contributions of the household head for each year he or she participated in the labor market. We converted all contributions to constant 2005 dollars. In computing social security contributions, we included both the contributions made by the worker and those contributed by the firm, and we accounted for different contribution rates applied to the self employed. Lastly, when computing the total accumulated amount of social security contributions, we assumed that past contributions earned a 2-percent rate of return.

Our analysis of the data reveals three key facts. First, young households collectively have made large contributions to social security. Table 1 shows the sum of the household head's total social security contributions made over his or her lifetime relative to the household head's average labor market income. Average labor income is computed by averaging the head's labor income earned in 1994, 1995, 1996, and 1998.⁹ The average (median) household head in our sample had made social security contributions equal to 85 percent (64 percent) of his or her average labor income.

Second, young households hold large amounts of unsecured debt.¹⁰ Specifically, Table 1 shows that 62% of our sample held some unsecured debt in 1999. Unconditionally, the average household in our sample had an unsecured debt-to-average income ratio of 36 percent. The average amount of household debt, conditional on having some unsecured debt, was equal to 58 percent of the head's average labor income.¹¹ Also, indebtedness is highly skewed in the full sample. Table 1 shows that ten percent of households with unsecured debt had debt-to-income ratios of more than 120 percent of the head's average income.

Lastly, the propensity to hold unsecured debt falls with age (results not shown). The median unsecured debt-to-average income ratio falls from 14 percent for households with heads between the ages of 25 and 27 to only 3 percent for households with

⁸See <http://www.ssa.gov/OACT/HOP/hopi.htm>.

⁹The PSID did not survey households in 1998, resulting in households' 1997 labor income to go unmeasured. Given the design of the PSID, income earned in year t is reported in survey year $t+1$.

¹⁰Gross and Souleles (2002), Durkin (2000), and Kennickell, Starr-McCluer, and Surette (2000), among others, have documented the increasing importance of unsecured debt — and, in particular, credit-card debt — in household portfolios.

¹¹Unlike the Survey of Consumer Finances (SCF), the PSID does not distinguish between total debt and revolving debt. Some households may have a positive amount of debt at any time during the month for transaction reasons. These households may intend to pay off the debt before they accrue any interest charges. We would like to focus on households that hold revolving debt (that is, households that carry balances forward from month to month). Of those households that report positive credit card balances in the PSID, most hold more debt than can be justified by a transaction motive. In our sample, we find that 46 percent of all households (or 75 percent of households with debt) have accumulated debt greater than one month's worth of income. This finding is consistent with data from the 1995 SCF, which finds that 56 percent of all households pay interest on their credit card balances on a monthly basis (Gross and Souleles 2002).

heads between the ages of 37 and 39. This suggests that changes in the social security program with respect to its effect on household unsecured debt holdings will have the greatest impact on households early in their lifecycle.

If the reason that households accumulate debt is liquidity constraints, we would expect that those households holding unsecured debt would also have a negative net asset position. To explore this question, we define three measures of household wealth. First, we examine household “total net worth,” which includes vehicle holdings, real estate equity (including main home equity), business equity, stocks, corporate bonds, cash, checking accounts, saving accounts, and Treasury less unsecured debt.¹² Second, we create a measure of liquid assets by subtracting illiquid assets from our measure of net worth. We define illiquid assets as business equity, housing equity, and vehicle equity. We refer to this measure as “net liquid assets.” Finally, we realize that cash, saving accounts, and checking accounts may be used to make monthly purchases. As a result, they are not measures of the stock of a household’s wealth. To account for this, we remove these resources from our measure of net illiquid assets. Essentially, our measure of “net financial wealth” is the sum of stock wealth plus bond wealth less non-collateralized debt.

Most of the households in the sample have positive net worth in spite of their unsecured debt. Appendix Table A1 shows that only 16 percent of the total sample and 26 percent of those with positive debt have negative net worth. However, most of total net worth is accounted for by home equity and assets that provide limited liquidity (net vehicle equity, for example). If we look at “net financial assets,” we get a completely different picture. Specifically, 82 percent of the households with unsecured debt have negative net financial assets, suggesting that households with non-collateralized debt are constrained.

Above, we show that the average young household carries substantial unsecured debt and has made substantial social security contributions. Our hypothesis is that the average young household would be made better off if it could forgo social security contributions while young and use the funds to reduce unsecured debt. This would require that those households that hold the unsecured debt be the same households that have made large social security contributions. The data bear out this claim. To measure the effect that social security contributions can have on debt accumulation, we construct “social security augmented debt” by subtracting the head’s accumulated social security contributions from total household debt. If the head’s accumulated social security contributions exceed total unsecured debt for a particular household,

¹²This is the full PSID wealth definition. See Hurst, Luoh, and Stafford (1998) for a thorough discussion. Up through the top 2 percentiles, the PSID wealth data compare very well with the SCF wealth data (Juster, Smith, and Stafford 1999).

then social security augmented debt is set to zero.

Table 1 shows what happens to the distribution of debt when we allow access to social security wealth. In the whole sample, the incidence of unsecured debt falls from 62 percent of households to 17 percent. Of those households holding unsecured debt, 72 percent can eliminate their debt by accessing their accumulated social security contributions. For those that cannot eliminate their debt, the level of indebtedness falls dramatically. As noted above, for households that hold unsecured debt, the 75th-percentile household has debt equal to 53 percent of the head’s average labor income; after gaining access to past social security contributions the 75th-percentile household has an unsecured debt-to-average-labor-income ratio of only 4 percent. Among those households with some debt, the debt-to-income ratio of the 75th-percentile household falls from more than half to just 4 percent. These numbers should be seen as upper bounds. Given the structure of the PSID, we can only measure the actual social security wealth of one member of the household. Debt, however, represents the debt of all household members.¹³ In summary, the PSID data show that those households holding large amounts of unsecured debt have simultaneously accrued large amounts of social security contributions.

In terms of aggregates, the decline in both total debt and total interest paid for the U.S. economy that would result from allowing households under 40 to access their past social security contributions to pay off their accumulated non-collateralized debt is large. The calculation is outlined in Table 2. Using PSID sample weights and census data, we calculate that allowing households to pay off non-collateralized debt with past social security contributions would reduce total household debt by more than 220 billion dollars and would reduce total annual household interest payments by 17.8 billion dollars per year — or more than 750 dollars for every household that has non-collateralized debt.

3 A life-cycle model

We now construct a model of life-cycle consumption with a social security program. Our household head enters the labor force at age 21, retires and earns no income after age 65, and dies at age 80 with certainty.¹⁴ We assume that labor is supplied

¹³We experimented with imputing the spouse’s social security wealth, given the spouse’s current age and work status. Under plausible assumptions, the percent of households with positive debt after allowing access to both spouses’ social security wealth falls by an additional 3 percentage points (from 17 percent to 14 percent).

¹⁴For our results on rule-of-thumb households, we must assume that households have an arbitrarily small but positive level of income in retirement.

inelastically and that household income evolves deterministically over the life cycle.

All households in our model value consumption streams using a time- and state-separable von Neumann-Morgenstern utility function such that:

$$U(\{c_a\}_{a=t}^T) = \mathbb{E}_t \left[\sum_{a=t}^T \beta^{a-t} \frac{\tilde{c}_a^{1-\gamma}}{1-\gamma} \right]$$

, where c_a is consumption at age a , \mathbb{E}_t is the expectations operator conditional on time- t information, and β is a time discount factor. As discussed above, we consider two types of households: Optimizing households that use financial assets to move consumption across states and dates and rule-of-thumb households that set consumption equal to income.¹⁵ We take no position on the reason that rule-of-thumb households behave the way they do. Some researchers, Reis (2004), for example, attribute such behavior to the large costs of acquiring and processing information. Alternatively, the combination of self-control problems, which limit savings, and credit restrictions, which limit borrowing for low-income households, would also yield similar behavior.

We allow optimizing households to trade three financial assets to maximize $U(\cdot)$: They can buy equities, which earn a stochastic net return \tilde{r}_E , bonds, which earn a net risk-free rate r_L , and they can engage in unsecured borrowing on which they pay the rate $r_B \geq r_L$. The period budget constraint is:

$$c_t - \theta_{B,t} + \theta_{L,t} + \theta_{E,t} \leq y_t - \theta_{B,t-1}(1 + r_B) + \theta_{L,t-1}(1 + r_L) + \theta_{E,t-1}(1 + r_E),$$

where $\theta_{.,t}$ measures investment or borrowing in the corresponding asset at time t . Households can buy unlimited positive quantities of stocks and bonds, but they cannot take short positions in equity or bonds, nor can they borrow negative amounts. Initial endowments of assets equal zero. We impose the condition that a household must pay off all debts before it dies, which implies that the household cannot borrow more than the present value of all its future labor income discounted at the borrowing rate r_B .

We solve for policy rules for optimizing households computationally, using methods developed in Judd, Kubler, and Schmedders (2002). For details on our numerical solution, see the appendix to Davis, Kubler, and Willen (2006). Table 3 lists key

¹⁵Recent work in the consumption literature motivates our decision to categorize consumers as either optimizing or rule-of-thumb. Researchers have shown that both optimizing households and rule-of-thumb households coexist in the data, particularly with respect to retirement. Bernheim, Skinner, and Weinberg (2001) argue that drops in household consumption at retirement “are consistent with ‘rule-of-thumb’... theories of wealth accumulation.” Both Hurst (2003) and Scholz, Seshadri, and Khitatrakun (2003) argue that what we call an optimizing model describes consumption behavior at retirement for roughly 80 percent of the population, while what we call a rule-of-thumb model describes the behavior of the remaining 20 percent.

features of our parameterization of the model. We draw the reader’s attention to several aspects of our parameterization.

First, our choice of asset returns follows Campbell (1999). We set the annual risk-free investment return to 2 percent, the expected return on equity to 8 percent, and the standard deviation of equity returns to 15 percent. All returns are in real terms.

Second, Davis, Kubler, and Willen (2006) present evidence that the borrowing rate exceeds the riskless lending rate by 10 percentage points. Of those 10 percentage points, charge offs account for 1.3 percent of the loan value. Conservatively, we assume that the marginal and average borrower are the same, and thus we specify a wedge equal to 8 percent, yielding a borrowing rate of 10 percent. Our model is in partial equilibrium, so we do not attempt to explain the origins of the wedge between the borrowing and lending rates. However, Dubey, Geanakoplos, and Shubik (2005) and Bisin and Gottardi (1999) present general equilibrium models in which the prices paid by buyers and the prices received by sellers of financial assets diverge because of asymmetric information.

Third, for the life-cycle income processes, we adopt parameter values estimated by Gourinchas and Parker (2002) from the Consumer Expenditure Survey (CEX) and the PSID, adjusted as described in Davis, Kubler, and Willen (2006). The Gourinchas and Parker (GP) labor-income series is after-tax, and we make the simplifying assumption that income taxation would be invariant to our proposed schemes.¹⁶ In constructing income profiles, GP strip out “categories of expenditure that do not provide current utility,” like medical care, and they hold family size constant over the life cycle.

3.1 Social security system

The social security system in our model works in the following way. Households pay a proportional tax while working. Upon retirement, optimizing households receive a lump-sum payment equal to the value of an annuity paying a fixed fraction of the average of the highest 35 years’ of income. Rule-of-thumb households receive the annuity.

Our system differs from the real U.S. social security system in three fundamental ways. First, both contributions and benefits are strictly proportional to income,

¹⁶Gokhale, Kotlikoff, and Neumann (2001) show that by redistributing income over the life cycle, pension plans can have adverse tax consequences. In some of our policy experiments, we potentially change the life-cycle profile of pre-tax income significantly (by increasing the employer contribution to social security); this could affect many aspects of household decision-making. We do not attempt to model these effects. However, we remind the reader that in our policy experiment of exempting households from social security contributions prior to age 30, these effects will be small.

whereas in the real social security system, contributions and benefits vary non-linearly with income. Focus on the contributions first. Under current U.S. law, social security contributions are a fixed portion of income up to a cap, meaning that household contributions to social security decline as a fraction of income as income rises. Such a contribution scheme contradicts the standard logic of life-cycle savings, that one should contribute more when one's income is relatively high. Contributions are strictly proportional to income in our model, so any calculation we do understates the burden of social security along this dimension.

In our model and in the real world, social security benefits depend on income in the highest 35 years' of income. However, in the real world, the proportion of average lifetime income that one receives in retirement depends on the level of one's lifetime income. Specifically, the marginal increase in retirement income for a dollar of labor income falls with rising income. Thus, the real social security system reduces relative income for those who have done relatively well and increases relative income for those who have done relatively poorly, providing a sort of income insurance for households. Since income is non-stochastic in our model, such a feature does not play a significant role, but it does mean that our model ignores a potential benefit of a social security system.

The second difference between our system and the real system is that we give a lump sum payment rather than an annuity to optimizing households in retirement. We do this because in our model an annuity reduces household welfare. To see why, note that in our model households can replicate an annuity with a lump sum if they so wish, and they can often do better, for example, if they want their consumption to slope down. This implies that, if we eliminate social security, welfare improves, even in the absence of any other distortion. We view this as a problem for the model in light of arguments, discussed in the introduction, that private annuity markets generally cannot replicate the annuity provided by social security. By eliminating the annuity feature from social security, we ensure that the method of provision of benefits is not a liability for social security. Our policy experiments typically have no effect on the retirement portion of social security – they guarantee exactly the same payment stream as the existing system. So whatever welfare benefits an annuity confers are preserved in our policy experiments.

Third, we build on an idea of Hubbard and Judd (1987), who propose that the social security system exempt young households from contributing. For example, we consider a social security system in which only households over the age of 30 contribute. To maintain the relationship between contributions and benefits, we adjust the level of contributions later in life so that the total discounted contributions

under all plans are identical. Below, we discuss this adjustment in depth.

Table 3 also shows the basic parameters of our social security system. We draw the reader’s attention to two possibly unfamiliar terms: “internal lending rate” and “internal borrowing rate.” The internal lending rate (ILR) is the implied rate of return on social security contributions in our system and is usually referred to as the “internal rate of return” in the literature. For example, in our model, the ILR equals the rate of return on investment of contributions that yields the retirement lump-sum benefit we select. We actually work backwards: We choose a level of contributions and an internal lending rate to get our lump-sum benefit. Leimer (1994) calculates that the internal lending rate on social security contributions of those currently entering the social security system is 1.7 percent. We round up and assume an ILR of 2.0 percent. In our baseline scenario, this implies that the replacement rate in retirement equals around 43 percent.¹⁷

The notion of an internal borrowing rate (IBR) relates to our analysis of alternative social security arrangements. In our simulations, we consider two policy experiments. First, we allow households to push back the point in the life cycle at which they start to contribute to social security. Second, we allow households to withdraw money from the social security system to pay off debts. To maintain balance between contributions and benefits, we increase the level of contributions later in the life cycle. We view the reduction in contributions (from either the exemption or the withdrawal) as a loan. In other words, households are allowed to increase their after-tax income today in exchange for a reduction in income in the future. Reducing contributions today at a cost of increased contributions later in life is tantamount to giving a household a loan from current social security contributions to be paid back in installments at some point in the future. We call the implicit interest rate on this loan the “internal borrowing rate.” In our baseline scenarios, we set the internal borrowing rate equal to the internal lending rate. We also consider scenarios in which the internal borrowing rate exceeds the internal lending rate. It should be noted that the existing social security system already allows some borrowing of this sort. For example, the earlier a household retires, the lower the level of benefits the household receives. In effect, when a household retires early, it borrows against future benefits. Simple calculations show that the implied internal borrowing rate roughly equals the internal lending rate when a household opts for early social security benefits.¹⁸

¹⁷The Social Security Administration stated in 2003 that the replacement rate had been stable at about 42 percent for the past 20 years (See http://www.ssa.gov/kc/fact_sheet_14_exp.htm).

¹⁸Authors’ calculation, using benefit and life expectancy information from the Social Security Administration web site.

3.2 Borrowing and social security

Whether households can borrow and on what terms plays a central role in the effectiveness of forced saving schemes. Essentially, if households can borrow, they can “undo” social security, and depending on the type of household, “undoing” social security can help or hinder policy goals.

For optimizing households, a forced saving scheme as defined above will either have no effect on welfare or make households worse off. If optimizing households can borrow and lend unlimited amounts at the riskless rate, then they can “undo” social security perfectly, in which case, social security has no effect on their lifetime budget constraint and thus no effect on lifetime welfare. However, if optimizing households face restrictions on borrowing, social security will typically make them worse off, because they can no longer undo it. Thus, for optimizing households, borrowing mitigates the costs of a forced saving scheme.

For rule-of-thumb households, if the households can borrow and lend unlimited amounts at the riskless rate, then they can “undo” the commitment aspects of social security perfectly. If households face restrictions on borrowing, then social security can make households with a taste for commitment better off, precisely because they cannot undo the commitment aspects of social security.

3.3 Effect of social security over the life cycle for optimizing households

In this subsection, we compare household behavior in the above model with household behavior in an otherwise similar model without a social security program.¹⁹ In the discussion that follows, we use the terms “social security wealth” to refer to the accrued amount of social security contributions and “private wealth” to refer to household stock and bond holdings less any non-collateralized debt. As discussed above, we assume that social security contributions earn a 2-percent rate of return.²⁰ We refer to the sum of private wealth and social security wealth as “total wealth.”

Our simulations show that for optimizing households, the existence of a social

¹⁹For a thorough discussion of the implications of this model with social security and how the model matches empirical facts about household consumption, portfolio composition, and non-collateralized debt holdings, see the extended discussion in Hurst and Willen (2004).

²⁰We understand that the term “social security wealth” is a bit of a misnomer in the sense that the households do not technically own the claim on their social security contributions in the same way that they own the claim on their equity holdings. The social security benefits that households will receive in retirement are subject to changes in government policy. However, given that we link current social security contributions with expected social security benefits, we refer to accumulated social security contributions as social security wealth.

security system leads to significantly higher levels of debt, lower private wealth, lower consumption, and lower utility than in a model with no social security. The deleterious effect of social security for optimizing households follows directly from our assumptions about borrowing. As noted above, a household that can borrow at the riskless rate can undo social security completely. For example, as Geanakoplos, Mitchell, and Zeldes (1999) point out, only a borrowing-constrained household stands to gain from investing social security wealth in equities, since an unconstrained household would have invested the optimal amount in equities anyway.

Figure 1 shows the results of our simulations with and without our baseline social security system. First, the upper left panel of Figure 1 shows that optimizing households that are forced to participate in social security accumulate significantly more debt than those who are not. Initially, the differences are quite small; at age 26, social security participation leads to only a small increase in debt. But by age 30, households that do not participate in the social security program stop accumulating additional debt, whereas debt accumulation by households participating in social security continues. The contrast is starkest at age 36. The households participating in the social security program have more than \$13,000 of debt (in 2005 dollars) while the otherwise-similar households without social security have none. Again, the reason that the households participating in the social security system accumulate more debt is that they need to counteract the forced saving imposed upon them by the social security program. Additionally, it should be noted that for optimizing households in the social security program, their social security wealth always exceeds the amount of their unsecured debt holdings (results not shown).

The lower left panel and the upper right panel of Figure 1 show, respectively, the amount of total wealth and the amount of private wealth held by optimizing households. As seen from the figures, optimizing households participating in the social security program have both higher total wealth and lower private wealth than those not participating in the social security program, up through the age of 60. By definition, the difference between the total wealth and private wealth is due to social security wealth.

It should also be noted that eliminating social security has two effects on household portfolio allocation. On one hand, young households in this model are liquidity constrained. The elimination of the social security program increases disposable income for such young households and, therefore, decreases their desire to borrow. In other words, removing the forced savings program of social security relaxes a binding liquidity constraint on younger households. However, at the same time, the elimination of the social security program enables households to invest in equities earlier in

their life cycle. Doing so enables them to increase their lifetime resources which, in turn, increases their desire to borrow when young. As shown in the upper left panel of Figure 1, young households still accumulate unsecured debt even after the social security program is removed.

In fact, in our model, households participating in the social security program do not start investing in equities until after the age of 40. After the age of 40, they allocate at most only 50 percent of their total wealth to equities. The reason for this is that social security wealth always remains at least 50 percent of their total wealth. However, if the social security system were completely eliminated, households would start accumulating private wealth after the age of 35. Moreover, households not facing the social security program allocate essentially all of their private wealth, when it is positive, to equities. Both of these effects result in non-participating households' earning a higher return on more of their savings over a longer portion of their lives.

The net effect of the increased debt while young and the lower allocation of savings to equities over the entire life cycle is depressed consumption throughout the entire life cycle for households that participate in the social security program. As shown in the lower right panel of Figure 1, the social-security-induced consumption gap is roughly constant at a little more than \$1000 until shortly before retirement. After retirement, the gap in consumption increases. Overall, participation in a social security program decreases consumption at every point in the life cycle for optimizing households that face a binding liquidity constraint.

Table 4 shows level of utility (in consumption-equivalent units) for households at the age of 21. We show the level of utility for both optimizing and rule-of-thumb households. Additionally, we show the population-weighted levels of consumption, debt, and total wealth for optimizing households and the population-weighted level of consumption for rule-of-thumb households. The population weighted levels are used to account for the fact that consumption, wealth, and debt change over the life cycle. The population weights, therefore, average the variables across households of differing ages.

In row 1 of Table 4, we show the results for our baseline model. The first thing to note is that the certainty-equivalent consumption of optimizing households is 7.7 percent higher than the certainty-equivalent consumption of rule-of-thumb households. Second, eliminating the social security program (row 2) yields welfare gains for optimizing households and large welfare losses for rule-of-thumb households. Specifically, certainty-equivalent consumption would increase by 7.7 percent for optimizing households (from 37.9 to 40.8) and would fall to zero for the rule-of-thumb households. Remember, the justification for the existence of the social security program in our

model was to sustain consumption for rule of thumb households in retirement. For rule-of-thumb households, the removal of the social security system causes consumption to fall to zero during their retirement years. Additionally, optimizing households not facing a social security system have higher consumption (by about \$3,300), lower debt (by about \$2,600), and higher private wealth (by about \$40,000) than otherwise-identical households that face a social security program.

Overall, participation in the social security program depresses consumption, increases debt, decreases saving, and decreases utility for optimizing households compared with their circumstances in a world with no mandatory social security.

4 Alternate social security policy proposals

The goal of this paper is to assess quantitatively proposals that would mitigate the utility loss to optimizing households of forced saving while at the same time preserving the commitment aspect of social security for rule-of-thumb households. In doing so, we consider two policy experiments. First, we evaluate a proposal to exempt completely young households from contributing to social security. We call this the *exemption* proposal. Additionally, we take households that are already in the social security system and have accumulated both social security wealth and unsecured debt, and we look at the benefits of allowing those households to borrow against their social security wealth to pay off debt. We call this the *payoff* proposal.

We evaluate both policy proposals using four criteria. First, does the policy proposal prevent optimizing households from simultaneously borrowing and holding social security wealth — the empirical fact that motivated this paper? More generally, what happens to household consumption, saving, and debt? Second, what are the welfare consequences for both optimizing and rule-of-thumb households? Third, we compare our proposals with much-discussed alternatives that seek to raise the internal rate of return, or in our parlance, the internal lending rate. And finally, for the exemption proposals, we ask how the proposal affects the average household net contribution to social security. For all our experiments, we require that the replacement rate in retirement equal or exceed the replacement rate in the existing social security system.

4.1 The exemption proposal

To begin, we propose changes in the age structure of social security contributions. We consider three age-based exemptions: no contributions before age 30; no contributions

before age 40; and no contributions before age 50. The latter two policies will require large changes in a household's marginal tax rate after the exemption expires. Such changes in the tax structure could induce changes in household behavior which we abstract from at this time. However, the before-age-30 exemption requires very small changes in the marginal tax rate after the age of 30. As a result, the labor supply effects of changing tax rates after the age-30 exemption expires would be negligible. Lastly, recall that we require the households to increase their contributions after the exemption expires to assure that their retirement replacement rate equals the replacement rate in the existing system.

For our analysis, we return to the four questions posed above. First, do these policy proposals prevent households from simultaneously borrowing and holding social security wealth? Yes — all three exemption programs ensure that households have no debt by the time they start contributing to social security. For the age-40 and age-50 exemptions, this result is not surprising. Figure 1 shows that even under the current social security system, households pay off most of their debt by age 40. But even with the age-30 exemption, households are out of debt by the time social security contributions begin. These results are shown in the top panel of Figure 2.

More broadly, the payoff proposal improves the general household balance sheet. Table 4 shows per capita levels of consumption, debt, and savings for optimizing households under the different exemption plans (Rows 3, 4, and 5 for the age-30, age-40, and age-50 proposals, respectively). As discussed above, row 1 of Table 4 shows the baseline model with the existing social security system, and Row 2 shows the results of the model with no social security system. The first thing to note is that the age-30 exemption leads to an almost complete elimination of unsecured debt. This fact is also shown in the top panel of Figure 2. Additionally, optimizing households increase consumption and accumulate more total wealth in response to the age-30 exemptions compared with optimizing households under the baseline social security system.

Why is it that both consumption and total wealth go up? An exemption, as we define it, can be interpreted as a loan in that it raises income today in return for a reduction in future income. There is, however, a key difference between an exemption and unsecured debt in our model. Instead of paying a credit card company 10 percent, an exempted household pays the government 2 percent. The delayed contributions to social security allow younger households to have higher disposable income. The reduction in debt service payments allows households to increase consumption throughout their life cycle. To maintain increased consumption once disposable income falls after the exemption expires, households accumulate more private wealth

when young.

Table 4 also shows that the longer the exemption period, the larger the increase in consumption, total wealth, and unsecured debt (comparing rows 4 and 5 with row 3). Why is it that a longer period of exemption leads to both higher total wealth and higher debt? The longer the exemption, the steeper the income profile when young. The steepness of the income profile induces more borrowing today. In other words, the exemption from social security while young does not completely remove the household's desire to borrow. The borrowing motive is exacerbated as the exemption length increases because the longer the exemption, the larger the increase in lifetime resources. Longer exemption periods result in higher lifetime resources because households are able to start investing in equities earlier in the life cycle. At the same time, the longer the exemption, the bigger the discontinuous drop in income when the exemption ends. The sharp future decline in income also causes a household to increase saving today.

Second, what happens to utility? The exemption proposals generate substantial increases in utility for both optimizing and rule-of-thumb households. The age-30 exemption increases certainty-equivalent consumption by roughly comparable amounts for optimizing and rule-of-thumb households: 3.4 and 3.3 percent, respectively. An age-40 exemption helps both types of households, but the gains are bigger for the optimizing households: 5.5 percent versus 4.8 percent. And an age-50 exemption generates a large welfare gain for optimizing households: 6.7 percent of annual consumption. In contrast, the age-50 exemption results in a small reduction in welfare for the rule-of-thumb household vis-à-vis the age-40 exemption. For the optimizing household, the age-50 exemption reduces the burden of social security by about 90 percent — for a 6.7 percent change in utility (row 5) vs. 7.5 percent change in utility from abolishing social security (row 2).

It is evident why optimizing households have welfare gains from the social security exemption policies. As discussed above, optimizing households with realistic income profiles are liquidity constrained while young. The social security program's mandatory contributions force young households to save when they would optimally wish to borrow (or save less). Exempting young households from social security, enables them to move closer to their first best life-cycle consumption profile while at the same time accumulating less unsecured debt.

One additional benefit of our proposal is that it also increases welfare for rule-of-thumb households. The reason is that rule-of-thumb households follow a sub-optimal consumption plan, given their preferences outlined in Section 3. The bottom left panel of Figure 2 shows the actual pre-social security contribution income profile, the

post-social security contribution income profile, and the optimal income profile, all for rule-of-thumb households. The optimal income profile is the income profile that would maximize the preferences outlined in Section 3, given that the household is choosing a consumption rule that sets consumption equal to disposable income. Notice, that the benefit to rule-of-thumb households of a social security program is the transfer they receive from the government that allows them to have positive consumption in retirement.

The reason that rule-of-thumb households are made better off from the exemption proposals is that the proposals move the post-social security income profile prior to retirement close to the optimal profile of income. This is illustrated in the bottom right panel of Figure 2. We want to stress that the exemption proposals do not change the income provided to rule-of-thumb households in retirement. In other words, we can maintain the existing rationale for social security while at the same time making both rule-of-thumb and optimizing households better off during their working years. The key reason for this is that both rule-of-thumb and optimizing households are made better off by a flattening of the disposable income profile.

Our third evaluation criterion was to compare our proposals to existing proposals currently being debated by policymakers. One common proposal is to allow households to allocate a portion of their social security contributions to equities. In terms of our model, this is analogous to increasing the ILR that households earn on their social security contributions. How do the utility gains from the exemption proposals compare with gains from simply raising the internal lending rate? For both optimizing and rule-of-thumb households, the answer depends on whether we change contribution rates when we increase the internal lending rate. Suppose we assume that contribution rates remain the same as they are currently. Then, for optimizing households, our exemption proposals clearly dominate the proposal of allowing households to invest part of their contributions in equities. Even an increase in the internal lending rate to 5 percent increases certainty-equivalent consumption by an amount less than that provided by the age-30 exemption. As seen in row 6 of Table 4, optimizing households that face the baseline social security system augmented by a 5-percent ILR (as opposed to a 2-percent ILR) increase their utility by only 2.3 percent. As discussed above, the age-30 exemption increases utility by 3.4 percent for optimizing households (row 3).

By contrast, for rule-of-thumb households, raising the internal lending rate is a winning proposition. This welfare gain for rule-of-thumb households should not be surprising. If we leave contribution rates constant and increase the ILR, replacement rates in retirement must increase dramatically. For example, a 5-percent ILR, with

contributions held constant, leads to a replacement rate of 88 percent of pre-retirement income. In that case, the decline in cash flow at the time of retirement for rule-of-thumb households is small. Stabilizing income during retirement has large welfare benefits for rule-of-thumb households.

However, if we simultaneously lower withholding rates as we increase the ILR so that the replacement rate in retirement remains constant at 43 percent, gains to rule-of-thumb households from raising the ILR shrink. For rule-of-thumb households, an increase of the ILR to 5 percent generates a welfare increase of the magnitude of those generated by the age-30 and the age-40 exemptions (see row 7 of Table 4). For optimizing households, however, the gains are much larger: An increase in the ILR to 5 percent yields welfare increases comparable to those provided by the age-40 exemption. Why are the effects so different? For rule-of-thumb households, the opportunity to invest at 5 percent is extremely valuable — since they don't invest in equity (or anything else). But for optimizing households, even a 5-percent ILR is still less than they can get by investing in equities at 8 percent.

Finally, how do our proposals affect the net household contributions to social security? As we said in the introduction, we make no assumption here about how social security is funded. As far as households are concerned, social security withholds contributions and pays out benefits. But social security is basically a pay-as-you-go system, and its long-run solvency depends on taking in as much in contributions as it pays out in benefits. In other words, the solvency depends on having non-negative net contributions.

By changing the internal borrowing rate, we can improve the solvency of social security while preserving some of the welfare benefits. Throughout, we have characterized exemptions as loans from the government to the household to pay the social security contributions each year. The exemption proposals discussed above implicitly assume an internal borrowing rate (the interest rate on the loans) equal to the internal lending rate that roughly equals the riskless lending rate. What if the government charged a higher borrowing rate? Since households already pay 10 percent for unsecured debt, why should we require the government to charge only 2 percent?

In rows 8 and 9 of Table 4, we show the effects of higher borrowing rates on households when we introduce an age-40 exemption and charge different internal borrowing rates. According to our results, social security could charge as much as 7 percent interest and still make both types of households better off. The main upside for social security is an obvious increase in net contributions to social security. Using a 5 percent borrowing rate, for example, would more than double net contributions using year 2000 population weights. Using 2020 population weights, a 5-percent

internal borrowing rate would turn a per-capita deficit for social security into a per-capita surplus. Even an age-30 exemption and an IBR of 5 percent results in utility gains for optimizing and rule-of-thumb households of 2.3 and 2.5 percent, respectively. In other words, simply by delaying the age at which social security contributions start, the government could require households to pay 5 percent more later in life (in present value terms) and the households would still be better off relative to the current situation.

4.2 The payoff proposal

According to Appendix Table A1, the average debtor household is around age 30 and has average debt around 60 percent of income. Suppose we take a hypothetical household that has debt and social security wealth equal to 60 percent of income at age 30. Consider a policy proposal that allows the hypothetical 30-year-old household to withdraw cash from social security to pay off its non-collateralized debt. To make sure this is an admissible plan, we require that the household increase its subsequent contributions from 10.6 percent of labor income to 12.8 percent of labor income, assuring us that the retirement replacement rate equals the 43 percent guaranteed by the existing system.²¹ This proposal is not relevant for rule-of-thumb households, since in our model such households do not accumulate debt. Row 1 of Table 5 shows the utility of optimizing households under the existing social security plan. We express the utility as the certainty-equivalent consumption of a household of age 30.²²

How well does this payoff proposal do in terms of increasing welfare for such an optimizing household? The payoff proposal does not eliminate the problem of simultaneous unsecured borrowing and social security investment. The top panel of Figure 3 shows debt and social security wealth with our proposal (“reduced debt”) and without (“existing”) going forward from age 30. Households use most but not all of their social security wealth to pay down debt. A small portion is used to finance increased consumption (bottom panel of Figure 3). To maintain the higher level of consumption, the household continues to borrow small amounts until age 33, at which point it starts to pay off the debt. Since households must continue to make social security contributions, our proposal does not eliminate the problem of simultaneous debt and borrowing, although it reduces it dramatically. As discussed above, eliminating the debt allows households to invest in equities earlier (and reduce

²¹12.8 percent would be the tax rate necessary to pay for the withdrawal of 60 percent of income from accumulated social security wealth at age 30 to pay off existing debt.

²²The results differ from those shown in Table 4 in that we are examining the change in utility for a hypothetical household that happened to have unsecured debt equal to social security wealth, equal to 60 percent of income at age 30.

debt payments) resulting in a wealth effect, and this in turn causes households to want higher consumption. This leads them to continue wanting to hold some unsecured debt after the age of 30.

Second, the payoff proposal greatly increases the utility of the optimizing households. Table 5 shows that our proposal (row 3) raises certainty-equivalent consumption by 2.1 percent for optimizing households versus the existing social security system (row 1). To put this in perspective, Row 2 of the table shows that eliminating social security altogether would raise certainty-equivalent consumption for optimizing households by almost 10 percent.

Third, the payoff proposal compares well with alternative proposals for social security reform. Rows 5 and 6 of Table 5 show the effects on utility of increasing the internal lending rate but otherwise keeping the social security system as is. As above, this is similar to the proposed policy of allowing households to invest some of their social security contributions in equities. Allowing households to use social security to pay off debt at age 30 yields a welfare payoff equivalent to an increase in the internal lending rate (ILR) of between 3 and 4 percent for a optimizing household.

To help illustrate the payoff proposal, we again suggest an alternative interpretation. What we are doing here is allowing households to spend money now (to pay off debt) on the condition that they pay money later in the form of increased social security contributions. In other words, our household is borrowing money. Specifically, our household borrows 60 percent of income and pays it back in 35 annual installments, each one equal to 1.7 percent of annual income — the difference between the standard social security contribution and the one necessary to pay off the debt. It is easy to see that the interest rate on the loan equals the ILR. But from whom is the household borrowing? Again, as we discussed above, the household is, in effect, borrowing from the government. By charging a higher IBR, the government could generate more revenue while simultaneously increasing the welfare of optimizing households. In Row 4 of Table 5, we show that even with a 5-percent IBR, the increase in utility would still be positive if we allowed 30-year-olds to use their social security contributions to pay off existing unsecured debt.

5 Conclusion and directions for future research

In this paper, we show that the current social security system leads many households to save at low interest rates and borrow at high interest rates. We show empirically that simply allowing households to use the money they have paid in to the social security system to pay off debt would allow many households to get out of debt completely

and others to dramatically reduce their exposure to high-interest unsecured debt. We then considered two policy experiments aimed at resolving this problem in the context of a life-cycle model with both optimizing and rule-of-thumb households. First, we considered options to change the age structure of social security contributions to prevent households from borrowing while they also contribute to social security. Particularly, we proposed exempting younger households from contributions altogether. Additionally, we considered allowing households to use the money in social security to get out of debt. With both of these programs, we required larger contributions later in life so as to leave the net present value of total contributions the same as under the current system. We found that both proposals, but, in particular, the former, solved the problem in question and led to significant increases in household welfare, consumption, and saving and to reductions in high-interest unsecured debt. Moreover, our exemption proposals led to increases in welfare for both optimizing and rule-of-thumb households.

Additionally, we showed that our options generated comparable and often higher welfare increases than popular proposals to increase the return on investment in social security. And they did so without any major administrative change to the social security system. There are no individual accounts. There is no uncertainty about returns. And it preserves the basic functions of social security in that it does not subject rule-of-thumb households to politically unacceptable risks during retirement.

With larger changes to social security, for example, the age-40 and-50 exemption, the assumptions that labor supply remains unchanged and that asset prices remain unchanged become increasingly untenable. But, as noted above, for one policy proposal — the age-30 exemption — neither extension should have a sizeable effect on our conclusions. First, the increase in withholding on labor income is very small — rising from 10.6 to 12.9 percent of income, comparable in magnitude to the increase in social security withholding that took place between 1980 and 1994. Second, Table 4 shows that an age-30 exemption has a small impact on consumption demand (a little more than a 1-percent increase) and a small impact on saving (a little more than a 4-percent increase). The main change is a dramatic reduction in consumer unsecured debt. Such a change could affect the wedge between borrowing and lending rates. Analyzing the general equilibrium effects of such a policy change on consumer credit markets would be a fruitful area for future research.

References

- [1] Abel, Andrew B. 1986. Capital accumulation and uncertain lifetimes with adverse selection. *Econometrica* 54: 1079-97.
- [2] Akerlof, George. 1998. Comment. *Brookings Papers on Economic Activity*. 185-189.
- [3] Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. 2001. What accounts for the variation in retirement wealth among U.S. households? *American Economic Review*, 91:832-857.
- [4] Bisin, Alberto, and Piero Gottardi. 1999. Competitive equilibria with asymmetric information. *Journal of Economic Theory* 87:1-48.
- [5] Bohn, Henning. 1999. Should the social security trust fund invest in equities? An intergenerational welfare analysis. *Review of Economic Dynamics* 2:666-697.
- [6] Campbell, John Y. 1999. Asset prices, consumption, and the business cycle. In *Handbook of Macroeconomics, Volume 1C*, edited by John B. Taylor and Michael Woodford. Amsterdam: North-Holland.
- [7] Campbell, John Y., Joao Cocco, Francisco Gomes and Pascal J. Maenhout. 2001. Investing retirement wealth: a lifecycle model. In *Risk Aspects of Investment-Based Social Security Reform*, edited by John Campbell and Martin Feldstein. Chicago, IL: University of Chicago Press.
- [8] Campbell, John Y., and N. Gregory Mankiw. 1989. Consumption, income, and interest rates: Reinterpreting the time series evidence. In *NBER Macroeconomics Annual*, edited by O.J. Blanchard and S. Fischer. Cambridge: MIT Press.
- [9] Davis, Steven J., Felix Kubler and Paul S. Willen. 2006. Borrowing costs and the demand for equity over the life-cycle. *Review of Economics and Statistics*, forthcoming.
- [10] Diamond, Peter A. and John Geanakoplos. 2003. Social security investment in equities. *American Economic Review* 93:1047-1074.
- [11] Dubey, Pradeep, John Geanakoplos and Martin Shubik. 2005. Default and punishment in general equilibrium. *Econometrica* 73(1):1-37.
- [12] Durkin, Thomas A. 2000. Credit cards: Use and consumer attitudes. *Federal Reserve Bulletin* (September):623-634.

- [13] Eckstein, Zvi, Martin Eichenbaum and Dan Peled. 1985. Uncertain lifetimes and the welfare enhancing properties of annuity markets and social security. *Journal of Public Economics* 26:303-26.
- [14] Feldstein, Martin, and Jeffrey B. Liebman. 2002. Social security. In *Handbook of Public Economics*, vol. 4, edited by A. Auerbach and M. Feldstein. Amsterdam: North-Holland.
- [15] Feldstein, Martin, and Elena Rangelova. 2001. Individual risk in an investment-based social security system. *American Economic Review* 91:1116-1125.
- [16] Geanakoplos, John, Olivia Mitchell, and Steven Zeldes. 1999. Social security money's worth. In *Prospects for Social Security Reform*, edited by Olivia S. Mitchell, Robert J. Myers, and Howard Young. Philadelphia: University of Pennsylvania Press.
- [17] Gokhale, Jagadeesh, Laurence J. Kotlikoff, and Todd Neumann. 2001. Does participating in a 401(k) raise your lifetime taxes? NBER Working paper 8341.
- [18] Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. Consumption over the life cycle. *Econometrica* 70:47-89.
- [19] Gross, David, and Nicholas Souleles. 2002. Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data. *Quarterly Journal of Economics* 117:149-85.
- [20] Hubbard, R. Glenn, and Kenneth L. Judd. 1987. Social security and individual welfare: Precautionary saving, borrowing constraints, and the payroll tax. *American Economic Review* 77:630-646.
- [21] Hurst, Erik. 2003. Grasshoppers, ants, and pre-retirement wealth: A test of permanent income consumers. University of Chicago working paper.
- [22] Hurst, Erik, Ming Ching Luoh, and Frank P. Stafford. 1998. Wealth Dynamics of American Families, 1984-1994. *Brookings Papers on Economic Activity* 1:267-337.
- [23] Hurst, Erik and Paul Willen. 2004. Social security and unsecured debt. NBER Working paper 10282.
- [24] Judd, Kenneth, Felix Kubler, and Karl Schmedders. 2002. Computational methods for dynamic equilibria with heterogeneous agents. In M. Dewatripont, L. P.

Hansen, and S. J. Turnovsky, editors, *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*. Cambridge: Cambridge University Press.

Juster, F. Thomas, James P. Smith, and Frank Stafford. 1999. The Measurement and Structure of Household Wealth. *Labour Economics* 6(2):253-275.

- [25] Kennickell, Arthur B., Martha Starr-McCluer, and Brian J. Surette. 2000. Recent changes in U.S. family finances: Results from the 1998 Survey of Consumer Finances. *Federal Reserve Bulletin* 86(1): 1-29.
- [26] Leimer, Dean. 1994. Cohort specific measures of lifetime net social security transfers. ORS Working Paper No 59, Office of Research and Statistics, Social Security Administration.
- [27] Reis, Ricardo. 2004. Inattentive Consumers. NBER Working Paper 10883.
- [28] Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2003. Are Americans saving optimally for retirement? University of Wisconsin - Madison working paper.
- [29] Smetters, Kent, and Jan Walliser. 2004. Opting out of Social Security. *Journal of Public Economics* 88(7-8):1295-1306.

Table 1: Debt, Wealth, and Social Security Contributions Relative to Average Income, by Debt Holdings

		I	II
		All	Have Positive Debt
Head's Total SS Contributions	Mean	0.85	0.81
	Median	0.64	0.62
Have Positive Unsecured Debt	Mean	0.62	1
	Median	0.36	0.58
Household Unsecured Debt	Median	0.05	0.22
	75th Percentile	0.29	0.53
	90th Percentile	0.84	1.23
Have Positive Social Security Augmented Unsecured Debt	Mean	0.17	0.28
	Median	0.19	0
Social Security Augmented Unsecured Debt	Median	0	0
	75th Percentile	0	0.04
	90th Percentile	0.33	0.74
Sample Size		2,077	1,233

Notes: Sample includes household heads in the 1999 PSID between the age of 22 and 39 who were either unmarried or married with a complete earnings history. Household unsecured debt is measured as of 1999. Column I shows descriptive statistics for our main analysis sample, while the sample in column II imposes the additional restriction that the household has positive unsecured debt. Head's total social security contributions are computed by the authors, using the household head's earnings history and social security tax formulas. The contributions are accumulated from the head's first year in the labor force through 1998 and are assumed to earn a 2-percent rate of return. See text for details. All statistics in this table (aside from whether or not the household has unsecured debt) are relative to the household head's average labor income. Average labor income is computed by averaging the head's labor income over the years 1994, 1995, 1996, and 1998. The PSID did not survey households about their 1997 labor income. Social security augmented debt is the amount of debt a household would have if it used all of the head's accumulated social security contributions to pay off its debt. By construction, social security augmented debt is non-negative.

Table 2: Aggregate Effects of Allowing Households to Use Social Security Wealth to Pay Off Non-collateralized Debt.

Average debt reduction for households with debt between the ages of 22 and 39 (1999)	\$9,579
Average annual interest saving resulting from debt reduction for households with debt between the ages of 22 and 39 (1999)	\$767
Total debt reduction for households aged 22 to 39 (1999)	\$221.8 billion
Total annual interest saved by households aged 22 to 39(1999)	\$17.8 billion

Notes: All dollar amounts are in 2005 dollars. Data come from the 1999 Panel Study of Income Dynamics. Debt reduction is calculated by using accumulated social security wealth to pay off existing non-collateralized debt. Non-collateralized debt assumes an interest rate of 10 percent while social security wealth assumes an interest rate of 2 percent. To compute total debt reduction for the U.S. population, we use 104 million households (Census Bureau), 36.6 percent of which have heads between the ages of 22 and 39 (PSID weighted data), and 60.8 percent of which have non-collateralized debt (PSID weighted data).

Table 3: Parameter Settings

Parameter	Baseline	Alternative values
Relative risk aversion	3	
Annual discount factor	0.97	
Age of labor force entry	21	
Age of retirement	65	
Age of death	80	
r_L	2%	
r_B	10%	
$E(\tilde{r}_E)$	8%	
$\text{std}(\tilde{r}_E)$	15%	
<i>Social security parameters</i>		
Start year	21	30,40,50
Contribution rate	10%	see tables
Internal lending rate	2%	see tables
Internal borrowing rate	2%	3%,4%,5%

Table 4: Households' Behavior Under Baseline Model and Various Policy Proposals.

#	Parameters of SS system					Utility		Population weighted <i>per capita</i>				
	Start age	Withholding	ILR	IBR	Repl. rate	Optimizing (% Δ)	Rule-of-thumb (% Δ)	Optimizing			Rule-of-thumb Consumption	Net SS contrib.
								Cons.	Debt	Total Wealth		
1	21	10.6	2	2	43	34.7 (0.0)	32.4 (0.0)	36.6	86.4	30.3	43.5	1.9
2	21	0.0	2	2	0	37.8 (9.1)	0.0 (-100.0)	39.7	63.3	55.2	45.9	0.0
3	30	12.9	2	2	43	39.2 (13.1)	36.1 (11.4)	45.1	0.3	103.8	43.1	2.3
4	40	17.9	2	2	43	40.0 (15.3)	36.6 (12.9)	45.8	1.4	123.6	43.1	2.3
5	50	30.4	2	2	43	40.4 (16.7)	36.6 (12.8)	46.7	2.2	145.0	43.7	1.7
6	21	10.6	5	2	88	38.8 (11.9)	38.1 (17.6)	46.5	10.3	66.5	45.8	-0.9
7	21	5.4	5	2	43	40.0 (15.5)	36.3 (12.0)	47.0	5.9	108.5	45.9	-0.5
8	40	24.0	2	5	43	39.2 (13.1)	36.0 (11.0)	44.6	0.7	129.8	41.3	4.1
9	40	30.3	2	7	43	38.4 (10.7)	35.2 (8.6)	43.5	0.3	138.3	39.5	5.9
10	30	14.7	2	5	43	38.8 (11.9)	35.8 (10.5)	44.3	0.3	100.9	42.4	3.0

Notes: Parameters of our baseline model are shown in row 1. “Start age” refers to the age at which social security contributions begin, while “withholding” refers to the rate at which households contribute to the social security program. “ILR” and “IBR” are the internal lending rate and the internal borrowing rate, respectively. See text for details. “Repl. rate” is the annuity value of the lump-sum payment from social security as a fraction of the top 35 years’ of income. “Utility” is the certainty-equivalent consumption that the household would receive starting at the age of 21, calculated under the model parameters. Changes in utility relative to the baseline model are shown in parentheses. Population-weighted parameters represent the average level of consumption, debt, saving, and net social security contributions in the model across the entire population. See the text for details. Rows 2 - 10 show the results of the different policy experiments discussed in the text. All dollar amounts are in thousands of 2005 dollars.

Table 5: Household Behavior under the Proposal of Allowing Households to Use Social Security Wealth to Pay Off Unsecured Debt

#	Wealth at Age 30		Parameters of SS system				Utility
	Private	SS	With- holding	ILR	IBR	Repl. rate	CE Cons. (% Δ)
1	-0.6	0.6	10.6	2	2	42	43.3 (0.0)
2	0.0	0.0	0.0	2	2	0	47.4 (9.5)
3	0.0	0.0	12.8	2	2	43	44.2 (2.1)
4	0.0	0.0	14.7	2	5	43	43.3 (0.1)
5	-0.6	0.6	10.6	3	2	51	43.7 (1.1)
6	-0.6	0.6	10.6	4	2	64	44.3 (2.4)

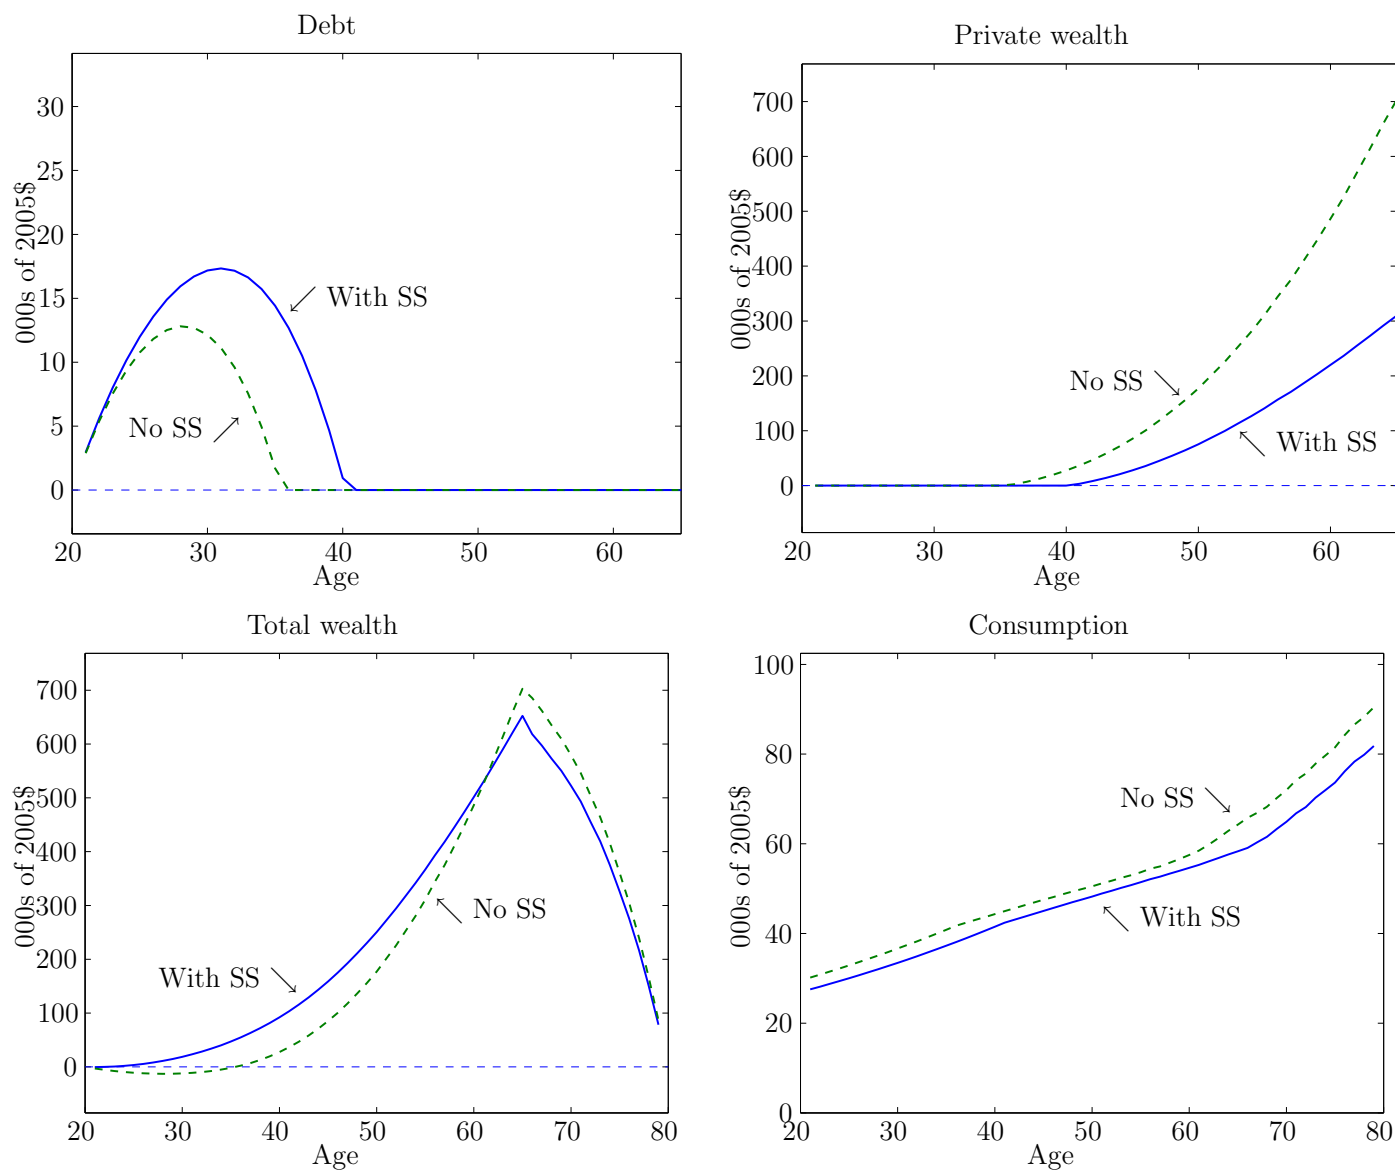
Notes: This table shows household behavior for a fictitious household that has private wealth at age 30 relative to permanent income (Private) of -0.6, and social security wealth at age 30 relative to permanent income (SS) of 0.6 such that total wealth at age 30 equals zero. “Withholding”, ILR, IBR, and “Repl. rate” are defined as in Table 4. Utility is defined as the certainty-equivalent consumption of this household at the age of 30. Row 1 shows the behavior of this household under the baseline model. Row 2 shows the behavior of this household under the baseline model with no social security program. Rows 3 and 4 show the behavior of this household allowing it to access its social security wealth at age 30 to pay off unsecured debts. Rows 5 and 6 show the results of the baseline model with higher ILR. See text for a full description of the policies discussed in each of the rows. All dollar amounts are in thousands of 2005 dollars.

Appendix Table A1: Descriptive Statistics for the 1999 PSID Sample, by Debt Holding.

		I	II	III	IV
		All	No Debt	Positive Debt	<i>p</i> -value of Diff
Age		31.6	32.3	31.1	<0.01
Dummy for Household Head=	Male	0.83	0.84	0.83	0.58
	Married	0.55	0.5	0.58	<0.01
	Black	0.02	0.16	0.09	<0.01
Average Labor Income	Mean	37,400	39,300	36,300	0.03
	Median	31,200	31,200	31,400	0.9
Dummy for Have Positive Cash		0.83	0.74	0.89	<0.01
Dummy for own	a Home	0.44	0.44	0.44	0.81
	Stocks	0.22	0.22	0.22	0.99
	Bonds	0.17	0.15	0.18	0.06
	Stocks or Bonds	0.32	0.3	0.33	0.26
	Stocks, Cash, or Bonds	0.85	0.77	0.9	<0.01
	a Business	0.09	0.1	0.09	0.91
Dummy for Negative	Net Worth	0.16	0.01	0.26	<0.01
	Net Liquid Assets	0.39	0	0.63	<0.01
	Net Financial Assets	0.5	0	0.82	<0.01
Net Worth to Average Income	Mean	2.02	3.07	1.38	<0.01
	Median	0.66	1	0.48	<0.01
Net Liquid Assets to Average Income	Mean	0.24	0.78	-0.08	<0.01
	Median	0	0.08	-0.08	<0.01
Net Financial Assets to Avg. Income	Mean	0.04	0.5	-0.24	<0.01
	Median	0	0	-0.16	<0.01
Sample Size		2,070	844	1,233	

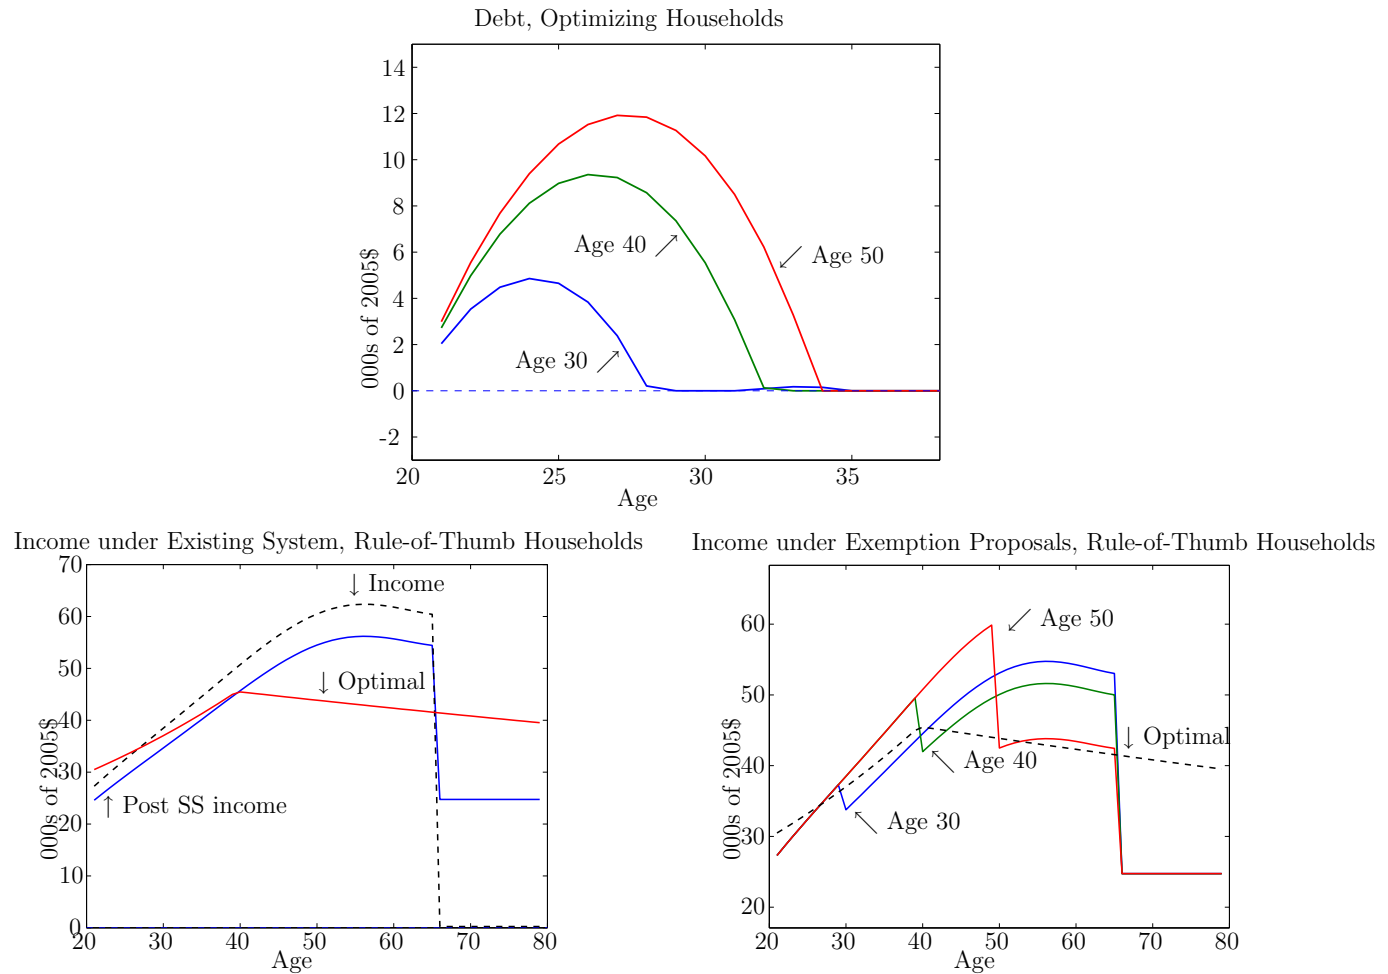
Notes: See note to Table 1 for full sample description. Average labor income refers to the average labor income earned by the household head between 1994 and 1998. See Table 1 for details. Household net liquid assets are defined as the sum of cash, checking, and saving balances, stocks, corporate bonds, and Treasury less unsecured debt. Household net financial assets are defined as the sum of stocks and corporate bonds less unsecured debt. All dollar amounts are in 2005 dollars.

Figure 1: Effect of Social Security on Various Household Indicators over the Life Cycle. Baseline Specification.



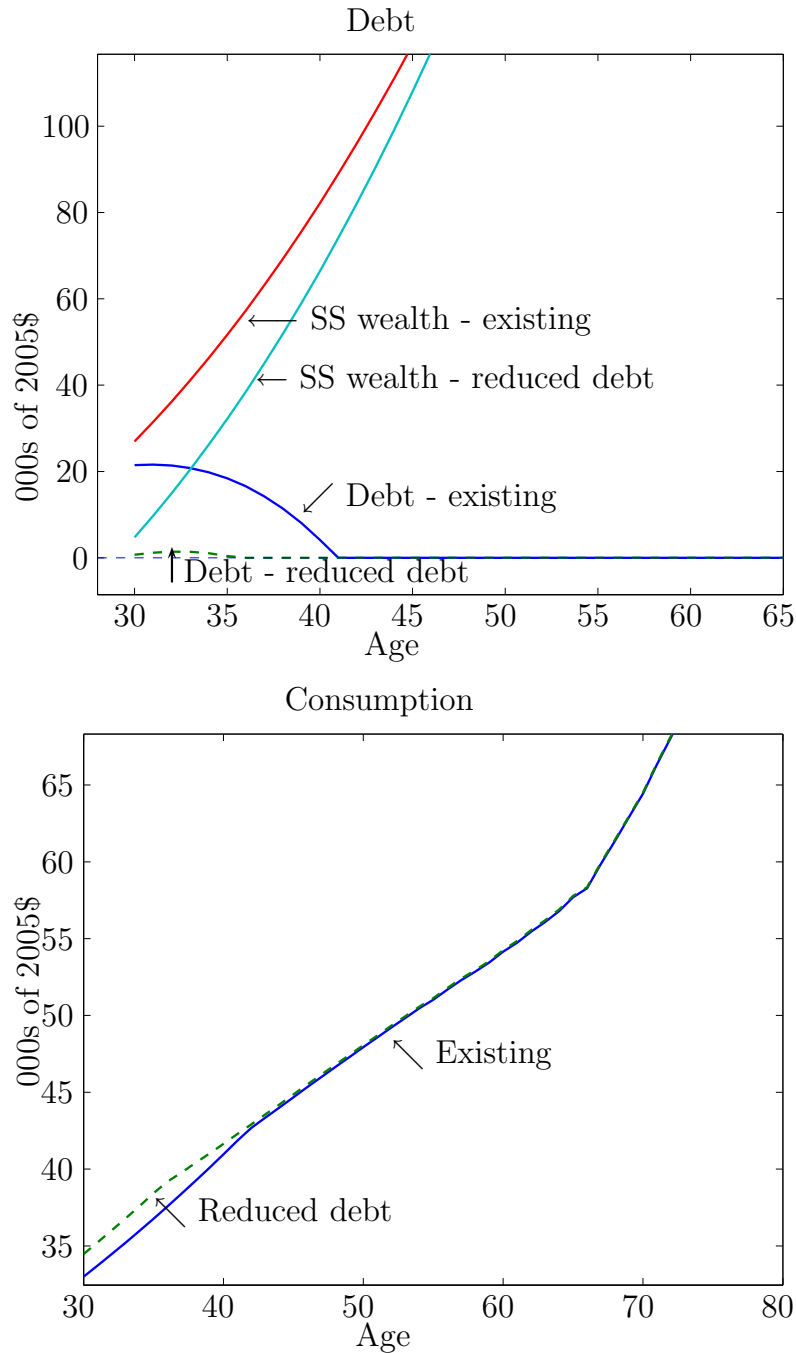
Notes: The figure displays the predicted unsecured debt, private wealth, total wealth, and consumption from our baseline model with social security (With SS), and our baseline model without social security (No SS). Total wealth includes both private wealth and social security wealth. See text for details. All amounts are shown in thousands of 2005 dollars.

Figure 2: The Effects of Exemptions over the Life Cycle.



Notes: The top panel of this figure shows the amount of debt holding for optimizing households in our model under our age-30 exemption proposal (Age 30), our age-40 exemption proposal (Age 40), and our age-50 exemption proposal (Age 50). The bottom left panel of this figure shows the actual income pre social security taxes (Income), the actual income post social security taxes (Post SS income), and the optimal level of income (Optimal) for rule-of-thumb households in our model. Optimal income would be the level of income that would maximize household utility for rule-of-thumb households, given that they set current consumption to current income. The bottom right panel of this figure shows the actual income post social security taxes for rule-of-thumb households under our three age-based exemption policies. All dollar amounts are in thousands of 2005 dollars.

Figure 3: Effects of Allowing Households to use Social Security Wealth to Pay Off Debt for a Household with Debt/Income Equal to 60 Percent.



Notes: This figure shows the effects of debt, social security wealth (SS wealth), and consumption under our baseline model (existing) and our debt-payoff proposal (reduced debt). See text for full details of our debt-payoff proposal. All dollar amounts are in thousands of 2005 dollars.