

SOCIAL SECURITY AND UNSECURED DEBT ^{*}

Erik Hurst[†]

Paul Willen[‡]

Graduate School of Business
University of Chicago
and NBER

Graduate School of Business
University of Chicago

July 2003

EXTREMELY PRELIMINARY AND INCOMPLETE.
DO NOT QUOTE WITHOUT AUTHORS' PERMISSION.

Abstract

Most young households simultaneously hold both unsecured debt on which they pay an average of 10 percent interest and social security wealth on which they earn less than 2 percent. We document this fact using data from the Panel Study of Income Dynamics. We then construct a dynamic life-cycle portfolio choice model and explore ways to reduce this inefficiency. We consider schemes that maintain the current level of retirement benefits. We show that both allowing households to use social security wealth to pay off debt and exempting young households from social security contributions (but in both cases requiring higher contributions later) leads to significant increases in consumption, saving and utility and reductions in debt. In particular, we show that exempting young households from social security contributions can dramatically reduce both the distortions caused by forcing young households to save in social security *and* the distortions caused by forcing them to invest all their social security wealth in low-yielding assets.

1 Introduction

The starting point for our analysis is the simple observation that households have about 700 billion dollars in non-collateralized debt on which the pay roughly 10

^{*}Last compiled: July 28, 2003

[†]Phone: (773) 834-4074. E-mail: erik.hurst@gsb.uchicago.edu.

[‡]Phone: (773) 834-5952. E-mail: paul.willen@gsb.uchicago.edu

percent interest and 11 trillion dollars of Social Security wealth on which they earn less than two percent interest.¹ As a nation, we are apparently doing exactly what we ridicule individual households for – borrowing on credit cards and saving with a passbook.

We focus on the old-age portion of social security and explore this topic in two steps. First, in Section 2, we tackle an obvious problem with our analysis. The above argument works for a representative agent, but in an economy with heterogeneous agents, it is possible that the households that have the debt and the households that have the wealth are different. We use data from the Panel Study of Income Dynamics to show that this is generally not the case. In our sample of households under the age of 40, 62 percent of households have non-collateralized debt. We show that if households could access their social security wealth to pay off debt, only 17 percent would still have debt. And our scheme would reduce debt significantly even for those households that still have debt: in our sample the 90th percentile household has non-collateralized debt equal to 84 percent of income. Our scheme lowers that number to 33 percent.

In Section 3, we construct a dynamic life-cycle portfolio choice model. We do this by adapting a model developed by Davis, Kubler and Willen (2003). In that model, households can invest in stocks and bonds and can also take out non-collateralized loans. We specify, realistically, that the interest rate on non-collateralized debt – the borrowing rate – exceeds the interest rate on bonds – the lending rate. We show that our parameterization can roughly match 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 of simultaneous debt and social security holdings. In our first experiment, we allow households currently in the social security system to access their social security wealth to pay off debt. In our second experiment, we build on an idea due to Hubbard and Judd (1987) and exempt young households from social security contributions. Both proposals lead to increases in consumption and saving, reductions in debt and substantial increases in lifetime certain equivalent consumption.

¹Our measure of non-collateralized debt is “consumer revolving credit” from the Federal Reserve Board – which was 725.0 billion dollars in May of 2003. Davis, Kubler and Willen (2003) argue that the interest rate on non-collateralized debt, taking default into account, is roughly ten percent in real terms. The value of social security wealth, defined as “the present actuarial value of the Social Security benefits to which the current population will be entitled at age 65 (or are already entitled to if they are older than 65) minus the present actuarial value of the Social Security taxes that they will pay before reaching that age,” comes from Feldstein and Rangelova (2001). Leimer (1994) calculates that the internal rate of return on social security contributions is 1.7%.

Let us elaborate briefly on our experiments. In our model, there is an optimal social security arrangement: none. Households can replicate social security privately if they want – but can do much better. We estimate for our baseline household that the welfare loss due to social security is around 7 percent of lifetime consumption. But we argue that the purpose of social security is to take care of households that aren't rational and forward-looking – i.e. households that don't behave like the households in our model. So we insist that any admissible proposal for social security must maintain the current level of retirement benefits. The goal of our experiments is thus to minimize what we call the portfolio loss – the loss to forward looking households due to suboptimal investment decisions of the social security system – subject to our admissibility constraint. So when we exempt a household from contributing to social security before age thirty, we also require that contributions after thirty increase by exactly enough to yield the same retirement benefit as the existing system. We assume no other changes to the social security system. In other words, households cannot invest social security contributions in equity; and we assume that the internal rate of return remains exactly the same as it is now.

Allowing households to withdraw money to pay down debt leads to significant reductions in the portfolio cost of social security. For example, we consider a household headed by a thirty years old 25th percentile of debt and social security wealth as measured in Section 2. Our plan reduces the portfolio cost for such a household from 8.4 percent of lifetime consumption to 7.2 percent of lifetime consumption. In other words, our plan raises lifetime consumption by 1.2 percent. This is equivalent to increase the internal rate of return on social security from 2 to more than 3 percent and entails no additional risk for either the household or the social security system.

How well do our exemptions do? Very well. As we mentioned above, the life-cycle portfolio loss of social security is about seven percent of lifetime consumption. Exempting households up to age 30 roughly halves the portfolio loss; exempting households up to age 40 cuts it by three quarters. And the portfolio loss when households don't need to contribute to social security until age 50 is less than one percent of lifetime consumption. Why are the effects so dramatic? It is easy to see that exemptions bring the household contribution schedule more into line with desired life-cycle saving. Most households under the age of thirty want to borrow not save. But what is surprising is that the exemption also reduces the distortions that result from the fact that social security requires investment of all contributions in an asset that pays two percent interest and does not allow any investment in equity which has historically returned 8 percent. We illustrate this surprising fact in Section 4.

Finally, we draw attention to one other possibility offered by our experiments. By

giving either an exemption or allowing a withdrawal but requiring higher contributions later, the government effectively loans money to households to maintain the value of the social security wealth. The conditions we impose imply that the government charges the internal rate of return on the loans. But the welfare gains from these “loans” are so big that the government could charge a higher interest rate and still make households better off. In other words, we show that the government could borrow at two percent, lend at five percent and still make households better off!

Before proceeding with the main body of the paper, we conduct a brief literature review. In our concluding section, we discuss some limitations of our analysis.

1.1 Literature Review

This paper builds on earlier work in two fields: social security and life-cycle portfolio choice.

Many researchers have explored the effects of social security on the economy. For a survey, see Feldstein and Liebman (2002). We draw the reader’s attention to two particularly relevant papers. Feldstein and Ranguelova (2001) propose that we replace the current pay-as-you-go social security system with individual accounts in which households can invest in equity and calculate the effects on retirement income. They find that people typically do better with individual accounts although there is a small probability that they will do worse. Hubbard and Judd (1987) look at a dynamic general equilibrium model of social security in which households face borrowing constraints. They show, among other things, that an exemption from social security contributions early in the life-cycle yields substantial welfare benefits.

Researchers have recently started to focus attention on the effects of borrowing limitations on life-cycle portfolio choice.² We build on this literature by extending a model due to Davis, Kubler and Willen (2003). Davis, Kubler and Willen limit borrowing not through a cap on the amount but by introducing a wedge between the cost of borrowing and the risk-free interest rate. They argue that a wedge is both empirically evident and that a model that incorporates a wedge leads to a more realistic model of life-cycle portfolio choice. Campbell, Cocco, Gomes and Maenhout (1999) explore the effects on portfolio choice and utility of allowing households to invest some of their social security wealth in equities in a life-cycle model with no borrowing. They find that a combination of smaller contributions and investment in equities leads to substantial welfare increases.

²Some examples include, Constantinides, Donaldson and Mehra (2002) who explore the effects of borrowing limitations on asset pricing. Gomes and Michaelides (2003) explore the effects of participation costs in conjunction with borrowing limitations.

Altig and Davis (1992) consider the effects of a wedge between borrowing and lending rates in an overlapping generations model. They find, as we do here, that the wedge dramatically affects the timing of bequests. Our findings, that a wedge affects the optimal timing of contributions, are similar in spirit.

2 Empirical facts about borrowing and social security

2.1 Data

The *PSID* is a large nationally representative survey started in 1968 which tracks the socio and economic variables of a given family over time. In each year of the survey, demographic questions such as age, race, family composition, and education levels are asked of all members of the households. Among other information, individuals report their labor market participation and earned labor income.

Occasionally, 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 their families holdings of cash, stocks, bonds, mutual 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 is the respondents' report of their holdings of non-collateralized debt (including store and credit card debt, student loans and other personal loans).

In this paper, we explore three measures of household wealth. For our main measure, we look at household "net financial saving" which is the sum of the household's holdings in stocks and corporate bonds less non-collateralized debt. We exclude cash, checking account, saving account, and Treasuries from this measure given that many households need a certain amount of liquidity on a monthly basis. However, to account for the fact that most households have large amounts of liquid reserves, we also examine household "net liquid saving" which adds cash, checking and saving account balances, and Treasuries to household net financial saving, defined above. Lastly, we examine household net worth which includes vehicle holdings, real estate equity (including main home equity) and business equity to household net liquid saving.

The *PSID* was designed to record repeated observations on the same household

across time. This makes the PSID ideal for measuring social security wealth. While the PSID does not have actual social security records for each survey respondent, social security wealth can be computed using the household's detailed earnings history coupled with the social security tax tables. The social security laws, as well as the exact procedure used to compute social security wealth, are defined in Appendix 1.

While 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. Given that the PSID was started in 1968, using the 1999 data allows us to follow an individual's labor market participation for upwards of 31 years. 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 non-collateralized debt. However, for completeness, we redid 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 given that the most non-collateralized borrowing occurs amongst younger households. This fact is consistent with the prediction of a lifecycle model given upward sloping earnings profiles. In any year of the PSID at least one member of each household either has been in the sample since the survey's inception (1968) or is a descendent of an original sample member. This implies that we have full earnings histories for all unmarried household heads between the ages of 22 and 39. However, for some married household heads, we will only have partial earning histories. The PSID classifies a household head as being male, if a male is present. If a man married into the PSID by marrying a woman PSID respondent, we would have only a partial earnings history for the male head and a full earnings history for his wife. Rather than imputing the expected previous earnings of some household heads, we restrict our analysis sample to only include married households where there are complete earnings histories for the head. In summary, our 1999 PSID sample includes all unmarried households with heads between the age of 22 and 39 and married households with heads between the age of 22 and 39 where the head also has a complete earnings history. The resulting sample size was 2,077 households. For robustness, we also examine only households where the head is unmarried. This sample included 850 households.

Table 1 shows descriptive statistics for our full analysis sample by propensity to hold non-collateralized debt. This table shows that households who have positive

non-collateralized debt tend to be slightly younger, more likely to be married and less likely to be black. Income of the household head is not a strong predictor, however, of the household's propensity to have debt. This result is not surprising given that both low income and high income households hold little, if any, non-collateralized debt. Low income households are unable to obtain debt, despite their desire to borrow, while high income households have less need to borrow. The median household with positive debt holds negative amounts of net financial assets and negative amounts of liquid assets. One of the most interesting facts from Table 1a is that the average indebtedness of households with non-collateralized debt is half the average total Social Security contributions (\$10,800 vs \$21,600).³

2.2 Discussion of data

Our analysis of the data reveals three key facts. First, young households have a lot of non-collateralized debt. Table 1 shows some descriptive statistics about our sample. 62 percent of households under the age of 40 have positive amounts of non-collateralized debt. The median debt to income ratio is five percent; the mean is 36 percent. Among just households that have some debt – i.e 62 percent of our sample – the median debt to income ratio is 22 percent and the mean is 58 percent.

Most of the households in the sample have positive net worth in spite of their non-collateralized debt. Table 2 shows that only 16 percent of the total sample and 26 percent of the positive debt sample have negative net worth. However, most household wealth is accounted for by home equity and assets which provide liquidity services. The return or shadow return of such assets may well exceed the interest rate on non-collateralized debt – so it's not clear that those funds should be used to pay down debt. So we define an alternative measure of net worth called “net financial assets” which excludes home equity and assets with liquidity services. Table 2 shows that 50 percent of the households in the sample have negative net financial assets and 82 percent of the households with non-collateralized debt have negative net financial assets.

Households generally see a small increase in indebtedness early in the life cycle as they accumulate debt. Table 5 shows that the median debt to income ratio rises from 13 percent between age 22 and 24 to 14 percent between ages 25 and 27 before falling to 3 percent between age 37 and 39. Variations for the 75th percentile household are more dramatic. Table 5 shows that the median debt to income ratio rises from 45

³all dollar amounts in the paper are reported as real 1996 dollars). The same result is found when looking at medians of non-collateralized debt and total Social Security contributions (\$5,700 vs. \$13,900, respectively)

percent between age 22 and 24 to 54 percent between ages 25 and 27 before falling to 20 percent between age 37 and 39.

Indebtedness is highly skewed. In the whole sample, the median household debt to income ratio is just five percent. Table 6 shows that a quarter have debt to income ratios of more than 29 percent and 10 percent have debt to income ratios of more than 84 percent. Just focusing on those households with some debt, we see that the debt is still quite skewed. 25 percent of households have debt to income ratios of less than 8 percent. But 25 percent have debt to income ratios of more than 53 percent. And 10 percent have debt to income ratios that exceed 120 percent.

Second, social security wealth is considerable even for young households. Results for both the with-debt and the without-debt samples are broadly similar. For the median household in our sample, the social security wealth to income ratio is 62 percent and for the mean household is 81 percent. Social security wealth rises rapidly over the life cycle. Table 5 shows that social security wealth to income ratio grows from 22 percent of income between 22 and 24 to 120 percent of income between age 37 and 39.

Third, the data bears out our basic claim that households could wipe out much of their non-collateralized debt if they had access to their social security wealth. Table 6 shows what happens to the distribution of debt when we allow access to social security wealth. In the whole sample, the incidence of non-collateralized debt falls from 62 percent to 17 percent. Of those households that have debt, 72 percent can eliminate it when they have access to social security wealth. For those that can't eliminate their debt, the level of indebtedness still falls dramatically. The 90th percentile household has debt equal to 84 percent income; after gaining access to social security wealth, only 33 percent are still in debt. 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.

3 A life-cycle model

We consider an optimizing model of household consumption and portfolio choice.

Our household enters the labor force and starts to receive labor income at age 21, retires and no longer receives labor income after age 65 and dies at age 80 with certainty. Our solution methods allow for stochastic variation in labor income but we choose to focus on the life-cycle aspects of the problem – and thus assume that labor income evolves deterministically. Evidence in Davis, Kubler and Willen (2003) shows that the effects of permanent shocks to income are similar to an increase in the

discount rate. Thus, one can think of our results as similar to those of a model with a lower discount rate and permanent shocks to income.

Households can trade three financial assets: They can buy equity with stochastic net return \tilde{r}_E , bonds at a net risk-free rate r_L , and borrow at the rate $r_B \geq r_L$. Households can buy unlimited positive quantities of stocks and bonds and can borrow unlimited positive amounts but they cannot take short positions in equity or bonds, nor can they borrow negative amounts. 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 .

A household chooses a plan for consumption, borrowings and asset holdings at date t to maximize

$$E_t \sum_{a=t}^T \beta^{a-t} \frac{\tilde{c}_a^{1-\gamma}}{1-\gamma}$$

where c_a is consumption at age a , E_t is the expectations operator conditional on time- t information, β is a time discount factor.

We solve the model computationally using methods developed in Judd, Kubler and Schmedders (1999). For details on our numerical solution, see the Appendix to Davis, Kubler and Willen.

Table 7 lists key features of our parameterization of the model. We draw the reader's attention to three aspects of our parameterization.

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

Second, Davis, Kubler and Willen (2003) present evidence that the borrowing rate exceeds the riskless lending rate by 8 percentage points. Thus we specify a borrowing rate of 10%.

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 Panel Study of Income Dynamics (PSID), adjusted as described in Davis, Kubler and Willen.

To carry out the welfare comparisons, we use certainty-equivalent consumption levels. To obtain certainty-equivalent consumption, we first calculate lifetime expected utility, U , for a given consumption profile. We then find the constant level of consumption, \bar{c} , that yields the same level of lifetime expected utility. That is, we

solve

$$\sum_{a=t}^T \beta^{a-t} \frac{\bar{c}^{1-\gamma}}{1-\gamma} = U \quad \text{for} \quad \bar{c} = \left[\frac{1-\gamma}{\sum_{t=0}^T \beta^t} U \right]^{\gamma-1}. \quad (1)$$

3.1 Social security system

Our Social Security system works in the following way. Households pay a proportional tax while working and receive a lump-sum payment on retirement. The lump-sum payment equals the value of an annuity paying a fixed fraction of the average of the highest 35 years of income.

Our system differs from the real Social Security system in three fundamental ways.

First, both contributions and benefits are strictly proportional to income whereas in the real Social Security system, contributions and benefits vary non-linearly with income. Let's focus on the contributions first. In the real world, Social Security contributions are a fixed portion of income up to a cap which has changed over time. Thus household contributions to social security are a declining fraction of income. In our model, there is no cap, so Social Security contribution are strictly proportional. This omission in the model actually strengthens our results. We argue below that the one problem with social security is that it forces households to save when their income is relatively low. The real system makes the problem worse by *lowering* the tax rate as income goes up. Now, let's talk about benefits. Social security benefits depend on income in the highest 35 years of income in both our model and in the real world. However, in the real world the proportion of your income that you receive in retirement depend on your income. Specifically, the marginal increase in retirement income for a dollar of labor income falls with income. To some extent, this provides a sort of income insurance for households, reducing relative income for those who have done relatively well and increasing relative income for those who have done poorly. Since income is non-stochastic in our model, such a feature would not play a significant role – but it does mean that our model ignores a potential benefit of a social security system.

Second, we give a lump-sum rather than an annuity in retirement. In the real world, Social Security provides a household with an annuity. In a world with uncertainty longevity, annuities are useful – insuring a household against unexpectedly long life. Hubbard and Judd (1987) argue that market failures in the private provision of annuities mean that this annuity feature of social security is welfare improving. But in our model, there is no longevity risk, and households can replicate the annuity with bonds if they want, so giving a household an annuity reduces welfare. In particular, the annuity feature of social security reduces welfare relative to a model in which peo-

ple provide for their retirement purely using private saving – even in the absence of any other distortions. So we avoid biasing our model against a social security system by eliminating the annuity feature from social security. We do however report the annuity value of the lump sum as a percentage of income for comparison purposes.

Third, we build on an idea due to 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 as described below.

Table 7 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 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. In our model, we assume a contribution rate of 10.0 percent⁴ Rather than choose a benefit rate, we choose an internal rate of return, which in turn implies a benefit rate. Leimer (1994), in a widely cited article, calculates that the internal lending rate on social security contributions of those currently entering the social security system is 1.7%. We round up and assume that the internal lending rate is 2.0%. In our baseline scenario, this implies that the replacement rate in retirement equals around 43 percent.

The notion of an “internal borrowing rate” relates to our analysis of alternative social security arrangements. In our simulations, we will consider two policy experiments. First, we will allow households to withdraw money from the social security system to pay off debts. Second, as mentioned above, we allow households to push back the point in the life cycle at which they start to contribute to social security. To maintain balance between contributions and benefits, we can do one of two things. First, we can reduce retirement benefits. In that case, we use the internal lending rate to calculate a reduced lump-sum payment at retirement. Second, we can increase the level of contributions later in the life-cycle. In this case, we can view the reduction in contributions (from either the exemption or the withdrawal) as a loan: we allow households to increase their after tax income today in exchange for a reduction in the future – which is tantamount to giving a household a loan to pay current social security contributions to be paid back in installments at some point in the future. We call the 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

⁴We could change this to 10.6 percent for accuracy.

rate.

3.2 Basic results of the model

In this section, we outline the basic features of life-cycle consumption and portfolio choice implied by the model. For a thorough discussion of the results of the model, see Davis, Kubler and Willen (2003). In the discussion that follows we will use the term “private saving” to measure total wealth of the household not including social security wealth. Note that our “lump-sum” assumption means that all wealth in retirement is “private saving.” Two aspects of the solution are relevant to our discussion of social security.

First, the relationship between equity holdings and the borrowing rate is non-monotonic. Consider a small modification to the baseline scenario in which the expected return on equity is 8 percent and the borrowing rate is 8 percent. With this specification, no one would ever borrow money to buy equity since such an investment would have zero expected return and would increase household exposure to risk. Suppose we lower the borrowing rate: Equity holdings go up. Why? Even a small reduction in the interest rate would turn borrowing to buy equity into a winning proposition. Thus, reductions in the borrowing rate increase equity demand. Suppose we raise the borrowing rate: Equity holdings also go up. Why? Households borrow for consumption purposes early in life and as they age, they pay off the debt. Until they pay off the debt, however, purchasing equity makes no sense – the expected return on borrowing to buy equity is now negative. Thus, if we raise the borrowing rate, households borrow less for consumption purposes which means they need to spend less time paying off their debts, which means that they can start buying equity sooner. So equity demand reaches a minimum when the borrowing rate equals the expected return on equity. Since we choose a borrowing rate of 10 percent, which exceeds, but not by much, the expected return on equity, we can say two things about our model. First, household demand for equity will be relatively low in all our simulations. And second, no household will ever simultaneously hold debt and equity.

Second, as one would expect in a life-cycle model, borrowing and thus saving and equity holding are highly sensitive to the shape of the age-income profile. Panels 1-5 of Table 8 show consumption, stock and bond holdings, debt and social security wealth for households age 22-39. Panel 1 shows our baseline specification in which the life-cycle profile is the “pooled” estimate from Gourinchas and Parker. Two features of the borrowing profile stand out. First, the level of borrowing is considerable – peaking at 45 percent of income between the ages of 28 and 30.. Second, for the baseline

household, borrowing has a hump-shaped profile. The household accumulates debt from age 21 until roughly age 30 and then starts paying it off. Panels 2-4 consider alternative specifications for the age-income profile. Panels 2 and 3 use estimates for high school and college educated households respectively. High-school educated households borrow significantly less than the pooled sample; college educated households borrow significantly more. Why? Figure 1 shows that the age income profile is less steep for high-school educated households, particularly before age 30. Figure 1 also shows that the age income profile for College educated households is much steeper than the pooled sample until age 30. One interesting thing to note is that college educated households only accumulate debt from age 22-24 – after that they decumulate. Panel 4 uses a flat income profile with the same average income as the baseline specification and illustrates the importance of the life-cycle profile. With a flat income profile, households never borrow and by the age 37-39 period, accumulate assets worth more than three times annual income; households in the baseline specification accumulate none. Panel 5 shows that increasing the subjective discount factor reduces debt and increases asset accumulation, as one would expect.

Third, with reasonable parameters, the model can generate life-cycle borrowing patterns similar to those we observe in the data. Our baseline household has a life-cycle debt accumulation pattern roughly matches that of the 25th percentile household in Table 5. The household in the model starts a little lower, accumulates more and its debt level peaks a little later. Both the high school educated household and the baseline household with a high discount factor generate borrowing patterns similar to the median household. The focus of this paper is on debtors, which motivates our choice of a household that is roughly the median debtor as our baseline.

Fourth, the gains from equity ownership are quite small in this model. How much is the right to trade equity worth to households? The column labeled “equity benefit” in Table 8 shows the percentage increase in certain equivalent lifetime consumption of a household which can trade equity compared to one that can’t. For our specifications with realistic age-income profiles, equity increases lifetime consumption by two to three percent. To put this number in perspective, the gains are on the order of 40 percent in the standard Merton-Samuelson version of this model. Why does equity ownership have so little effect on lifetime consumption? The reason is simple. In the Merton-Samuelson model, households can borrow at the return on riskless bond – 2 percent here. Households can realize arbitrarily large returns with no investment by borrowing at two percent and investing in equity with an expected return of eight percent. Crucially, they can also use low interest unsecured borrowing to borrow against these excess returns and increase their current consumption accordingly. In

our model, equity does afford substantial increases in consumption; however, those increases occur late in life and thus have a comparatively small impact on lifetime utility.

3.3 Effect of social security over the life-cycle

Our simulations show that social security leads to significantly higher levels of debt, lower savings, lower total asset accumulation, and lower consumption throughout the life-cycle.

We first consider our baseline household through the life cycle. Panel 1 of Table 9 shows results of our simulations with our baseline social security system. Panel 2 considers the same household but without any social security at all. What are the differences?

First, the household with social security accumulates significantly more debt. Initially, the differences are quite small – at age 26, social security leads to only a one thousand dollar increase in debt. But at that point, the household without social security stops accumulating additional debt. By age 34, the no social security household has reduced its debt load to about three thousand dollars, while the household with social security has accumulated more than ten thousand dollars in debt. Why don't households stop borrowing completely when we eliminate social security? Eliminating social security has two effects. On one hand, it increases current income relative to retirement income – which increases saving. At the same time, households are liberated from the internal rate of return for a large portion of their investments. Figure 3 illustrates this. The line marked “existing” shows the percentage of household saving (including social security wealth) allocated to equities in the existing system – which is zero until age forty and never exceeds fifty percent. The line marked “optimal” shows the percentage for a household which faces no social security system – which equals one hundred percent except for a couple of years shortly before retirement. The change in portfolio returns tilts household income and induces borrowing early in life.

Second, the household with social security loses money on its total portfolio until age 38. Until that point, interest payments on debt exceed effective payoffs on social security investment. Third, consumption is higher throughout the life-cycle for the household without social security. The gap is about 1500 dollars at age 22 and widens to about 4000 dollars in retirement (at age 66). Overall, social security reduces lifetime utility by a little more than seven percent.

We now consider one of the households in our sample. According to our data, the 75th percentile household at age 30 has debt equal to approximately 40 percent of

income and social security wealth equal to 60 percent of income. Panel 3 of Table 9 traces out the evolution of various quantities for such a household according to our model. As in the data, this household only slowly pays off debt. By age 38, our model predicts that our household will have debt equal to approximately 10 percent of income – in the data it is more than 20 percent of income. What would happen if we eliminated social security, and paid back the accumulated contributions? Panel 4 shows what would happen. The household receives a check for 15,800 dollars from Social Security. What do they do with the money? First, the household pays off 8,500 of debt. Second, the household increases consumption by 3,300 dollars. And they invest the remaining 4,000 dollars in stocks. By retirement, the household that withdraws from social security has 13 percent higher net worth and 13 percent higher consumption. Leaving the social security system increases certain equivalent wealth by 8.4 percent.

4 Policy proposals to fix social security

We consider two policy experiments. First, 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 using social security wealth to pay off debt. Second, we consider the whole life-cycle and try to design a social security system that doesn't force households to both borrow and save at the same time.

Before we consider the examples, let us lay out our analytical framework. In our model, it is easy to see what the optimal social security scheme is: eliminate the program altogether. Since households can replicate the saving and investment program of the social security system, they can do at least as well on their own as they do with the system.⁵ So why have social security? Clearly, a rational, forward-looking household can do better than social security but anecdotal and empirical evidence suggest that some people aren't rational and forward-looking and don't plan ahead. For them, social security is a lifesaver, forcing them to provide for their retirement. So the starting point of our analysis here is that providing the current level of social security benefit in retirement is a condition of any admissible plan for social security. Such a condition necessarily imposes a cost on forward looking households – because it may force them to save more than they want to and because it forces them to

⁵As mentioned above, we abstract from longevity risk here. With longevity risk, social security would confer some benefits. But our focus here is on the accumulation phase of social security. If the problem is market failure of annuity markets, as suggested by Hubbard and Judd, we could remedy that without a forced saving scheme.

invest in a riskless bond paying two percent interest. We measure this cost as the percentage reduction in certain equivalent consumption from the no social security benchmark and we call it the *portfolio cost* of social security. The goal of our policy proposals is to minimize the portfolio cost of social security subject to the constraint that we guarantee the current level of social security benefits in retirement.

To understand the portfolio cost, recall our examples in Table 9. In Panels 1 and 2 we compared consumption over the life-cycle with and without social security. The portfolio cost of social security is $(23.9 - 22.2)/23.9 = 7.0$ percent. In Panels 3 and 4, we compared consumption from age 30 on for a household with offsetting positions in debt and social security. The portfolio cost of social security in this example is $(28.1 - 25.7)/28.1 = 8.4$ percent. Note that going from Panel 1 to 2 or Panel 3 to 4 are not admissible policy changes. In both cases, the replacement rate in retirement is zero percent – we have shredded the safety net!

We also construct another measure of the gain of our proposal. Many people have argued that the problem with social security is that the internal lending rate (ILR) is too low – investors should be able to do better than two percent. So to compare our proposal to broad attempts to increase the ILR, we consider the following experiment. Suppose we made no other change to social security other than raising the ILR. By how much would we have to raise the ILR to reduce the portfolio cost by as much as our proposal? We call this quantity the *equivalent return*.

4.1 Buying out debt

Suppose we have a 30 year-old household with debt equal to 40 percent of income and social security wealth equal to 60 percent of income (just as in Panel 3 of Table 9). Consider the following policy proposal. Allow the household to withdraw exactly enough to pay off its debts. To make sure this is an admissible plan, we require that the household increase its subsequent contributions from 10.0 percent of labor income to 11.7 percent of labor income which assures us that the retirement replacement rate equals the 43 percent guaranteed by the existing system.

How well does our policy experiment do? Panel 1 of Table 10 shows utility, consumption and total wealth, under the existing system which repeats information from Panel 3 of Table 9). Panel 2 shows what would happen if we followed our proposal. We note three things. First, the household does not pay off all its debt. Compare that with Panel 6 (which repeats information from Panel 4 of Table 9) – in which we shut down social security altogether and the household does pay off all its debt. The difference here is that our proposal only gives the household enough to pay off its debt. Thus to increase consumption, the household must retain some debt.

Second, consumption increases throughout the life-cycle but the change dissipates to almost zero by retirement. Increases in non-social security wealth lead to increased consumption while working, not in retirement. Third, Our proposal reduces the portfolio loss of social security from 8.4 percent of lifetime consumption to 7.2 percent. The increase in certain equivalent consumption equals roughly 400 dollars per year. Fourth, our proposal is equivalent to raising the internal lending rate to a little more than 3 percent. Panels 3-5 show the effects of raising the ILR and leaving everything else the same. Raising the ILR *increases* debt – as households anticipate higher income in retirement. Private saving falls as well. In our proposal, by contrast, private saving increases and debt falls.

To help illustrate the virtues of our proposal, we suggest an alternative interpretation. What we are doing here is allowing household to spend money now (to pay off debt) on 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 40 percent of income and pays it back in thirty-five 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 who is the household borrowing from? One way to view it is that the household is borrowing from itself. Later on, we will argue that the sensible interpretation is as a loan from the government but when the internal borrowing and lending rates are the same, either interpretation is valid. So what the household is really doing is investing social security wealth in a loan to itself. Our equivalent return calculations show that reallocating a social security portfolio into loans to oneself is equivalent to reallocating your social security portfolio into an asset that pays a little more than 3 percent return for sure.

4.2 Changing the life-cycle structure of social security

In this section, we propose changes in the age structure of social security. Specifically, we build on Hubbard and Judd’s idea (1987) to ameliorate the life-cycle effects of social security by providing young households with an “exemption” from social security contributions. To make this an admissible proposal, we still require the households increase their contribution after the exemption expires to assure that their retirement replacement rate equals the replacement rate in the existing system. We consider three age-based exemptions: no contributions before age 30; no contributions before age 40; and no contributions before age 50.

Table 11 shows that the effects of our proposal on our baseline household are

dramatic. The portfolio loss of the baseline scenario is 7.0 percent. An exemption up to age 30 nearly halves the portfolio loss, cutting it to only 3.8 percent. An exemption up to age 40 reduces portfolio loss by almost three quarters to 1.9 percent. An exemption up to age 50 reduces portfolio loss by almost 90 percent to 0.8 percent.

Why do these exemptions have so much power? We now show that our exemptions are successively better and better approximations to the optimal saving and portfolio choice decisions. We first consider the saving decision and then the portfolio allocation decision. Figure 2 shows total saving (including social security wealth) in excess of the no-social security optimum for a baseline household under four scenarios: the existing social security system; and social security with the three age-based exemption rules. As the picture shows, deviations from the optimum get progressively smaller as we extend the exemption.

We note two odd features of the Figure 2. First, households in all the social security systems oversave relative to the no social security benchmark until roughly age sixty. For households in the existing social security system, this is not too puzzling – the system forces them to save when they would actually like to borrow. But in the age 40 and age 50 exemptions, forced saving starts *after* the household would have started saving anyway. So why do they oversave even before the exemption starts. The reason is that the exemption leads to a discontinuous but predictable drop in income (17.3 percent for the age 40 exemption and 29.8 percent for the age 50 exemption, see Table 11.) Households engage in private saving in anticipation of this – which you can see in Figure 4, which shows what we call excess private or non-social security saving.

The second odd feature of the life-cycle excess saving profile is that households in all the social security systems undersave relative to the no social security benchmark after roughly age 60. Why does this happen? Figure 4 shows that private saving falls relative to the non-social security benchmark when the exemption ends – the forced saving crowds out private saving. Since the return on private saving is much higher than the return on forced saving, wealth accumulates much more quickly for households that don't have social security.

We now discuss how social security affects portfolio allocation and show that exemptions reduce the distortions generated by social security. Figure 3 compares portfolio allocations to equities over the life cycle for the four social security schemes and the no social security optimum. In the absence of social security, households invest all their money in equities until quite late in the life cycle – which we see in the line marked “Optimal” in Figure 3. Households invest almost all their money in equity in virtually all portfolio choice models with labor income which follows from

the fact that financial wealth is a small fraction of total wealth – which includes human capital – until late in the life cycle. Thus households can invest almost all their financial wealth in equity but still have relatively low exposure to it. Under the existing social security system, equity’s share never exceeds fifty percent and is zero for much of the life-cycle. Before the age of 40, households facing the existing social security system engage in no private saving – all their saving is forced and all forced saving is invested at the internal lending rate – 2 % in our case.

Figure 3 shows that exemptions close the gap. For the age 40 and 50 exemptions, households never invest less than 50 percent in equities and invest 100 for much of the life-cycle. Why? The basic reason is that with exemptions, households engage in considerably more private saving. Consider first a household with the age 30 exemption. They start saving before they reach 30 in anticipation of a decline in income when the exemption ends. At age 30, they start dissaving to maintain consumption, but by the age of 35, they are accumulating again and all their accumulation goes into equity. Now consider a household facing an age 40 or age 50 exemption. In both cases, they save entirely in equities in anticipation of the exemption. Then when the exemption ends, they start to draw down their private saving and replace it with forced saving, which they cannot invest in equities.

To quantify the effect of exemptions on optimal saving and on optimal portfolio allocation, we conduct the following exercise. Table 12 shows the effects of exemptions in the Hubbard-Judd scenario in which there are no risky assets. We note three things. First, the increase in certain equivalent consumption from eliminating social security altogether is much smaller – roughly 1000 dollars compared with the 1700 in the risky asset scenario. Second, an age 30 exemption wipes out about 70 percent of the portfolio loss and an age 40 exemption wipes out 97 percent. By contrast, when there are risky assets, an age 30 exemption only eliminates 45 percent of the portfolio loss and an age 40 exemption eliminates about 70 percent. As figure 3 shows, going from the age 40 exemption to the age 50 exemption brings portfolio allocation much closer to the optimal level.

Finally, what is the equivalent return of these proposals? Looking at Panels 5-7 of Table 11 shows that the equivalent return for all proposals exceeds 5 percent. In other words, even the age 30 exemption is better than raising the internal lending rate on social security to 5 percent with no risk. If we really could raise the internal lending rate without any additional risk, then an alternative admissible policy would be a combination of a reduced contribution rate and an increase ILR. Panels 8-10 show the minimum contribution rate that guarantees the existing retirement benefit for 3 percent, 4 percent and 5 percent ILR’s respectively. The exemptions still do

surprisingly well. The age 30 exemption matches the four percent ILR; the age 40 exemption matches the five percent ILR; and age 50 beats all three.

4.3 Raising the internal borrowing rate

One purpose of social security reform is to reduce the “portfolio cost” as we have defined it above. But another possible goal is to reduce the unfunded liability of the system. We now suggest that exemptions could be used to reduce the unfunded liability as well.

In Section 4.1, we characterized exemptions as loans from the government to the household to pay the social security contribution that year. The exemptions, as discussed in the previous section, assume an internal borrowing rate – the interest rate on the loans – equal to the internal lending rate which roughly equals the riskless lending rate. What if the government charged a higher borrowing rate? Since households already pay 10 percent for unsecured debt as it is, why should we require the government to charge only two percent? In Table 13, we show the effects of higher borrowing rates on households when we introduce an age 40 exemption and charge different internal borrowing rates. What we see is that even with an internal borrowing rate of 5 percent, the age 40 exemption still reduces the portfolio loss from social security by 2.3 percent. In other words, the government could borrow at 2 percent, lend at 5 percent *and* make households better off!

5 Conclusion and directions for future research (incomplete)

There are two limitations of the model which potentially affect our results. First, we assume exogenous labor supply. Our policy proposals introduce major changes in both the level of social security taxes and the life-cycle profile of them. For example, compared to 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 tax code could lead to major changes in life-cycle labor supply. However, there is a potential remedy. In keeping with our notion of an exemption as a loan, we could explicitly keep track of missed contributions before 50 and insist that the household repay them with interest after 50 in full. In other words, while exempt, households would accrue a liability proportional to their current income – so there would be a marginal tax on their income. After the exemption expires, households would pay the usual 10 percent social security contribution proportional to their income after the exemption

expires plus a fixed payment related to the value of the loan they took. Thus the marginal tax rate after the exemption ends would stay exactly as it is now.

Second, we have constructed a partial equilibrium model. Our policy proposals would lead to significant changes in equity demand, credit demand and the life-cycle consumption profile. It is easy to imagine that such changes would lead to changes in the overall price level that might erode some of the benefits we find here.

References

- [1] Altig, David and Steven J. Davis, 1992, The timing of intergenerational transfers, tax policy and aggregate saving, *American Economic Review*, 82(5), 1199-1220.
- [2] Campbell, John Y, 1999, Asset prices, consumption and the business cycle, in John B. Taylor and Michael Woodford, editors, *Handbook of Macroeconomics, Volume 1C*. Amsterdam: North-Holland.
- [3] Campbell, John Y. and Luis Viceira, 2002, *Strategic Asset Allocation*. New York: Oxford University Press.
- [4] 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*, John Campbell and Martin Feldstein, eds. Chicago, IL: University of Chicago Press.
- [5] Constantinides, George M., John B. Donaldson and Rajnish Mehra, 2002, Junior can't borrow: A new perspective on the equity premium puzzle, *Quarterly Journal of Economics*, 117(1), 269-296.
- [6] Davis Steven J. Felix Kubler and Paul S. Willen, 2002, Borrowing costs and the demand for equity over the life-cycle. NBER working paper XXXX.
- [7] Feldstein, Martin and Jeffrey B. Liebman. 2002. Social security. In A. Auerbach and M. Feldstein, editors, *Handbook of Public Economics*, vol. 4. North-Holland.
- [8] Feldstein, Martin and Elena Rangelova. 2001. Individual risk in an investment-based social security system. *American Economic Review* 91(4):1116-1125.
- [9] Gomes, Francisco and Alex Michaelides, 2002, Life cycle asset allocation: A model with borrowing constraints, uninsurable labor income risk and stock-market participation costs, working paper, London Business School.
- [10] Gourinchas Pierre-Olivier and Jonathan A. Parker, 2002, Consumption over the life cycle, *Econometrica*, 70(1), 47-89.
- [11] 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(4): 630-646.

- [12] Judd, Kenneth, Felix Kubler and Karl Schmedders, 2002, Computational methods for dynamic equilibria with heterogeneous agents. Forthcoming in M. Dewatripont, L. P. Hansen and S. J. Turnovsky, editors, *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*. Cambridge University Press.
- [13] Leimer, Dean. Cohort Specific Measures of Lifetime Net Social Security Transfers. ORS Working Paper No 59, Office of Research and Statistics, Social Security Administration, February 1994.
- [14] Tauchen, George and Robert Hussey, 1991, Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica*, 59(2), 371-96.

Table 1: Descriptive Statistics for 1999 PSID Sample, by Debt Holding. Notes – Data is from the 1999 PSID. Sample includes 1) all unmarried households with heads between the age of 22 and 39 and 2) married households with heads between the age of 22 and 39 where the head also has a complete earnings history. Head’s Total Social Security contributions are computed by the authors using the household heads earnings history and Social Security tax formulas. The contributions are accumulated from the head’s first year in the labor force through 1998. Household net liquid assets are defined as the sum of cash, checking and saving balances, stocks, corporate bonds, and Treasuries less non-collateralized debt. Household net financial asses are defined as the some of stocks and corporate bonds less non-collateralized debt. Sample is split by households with no non-collateralized debt (column II) and households with some positive amount of non-collateralized debt. All dollar amounts are reported in 1996 dollars. [tabk1]

		<i>I</i>	<i>II</i>	<i>III</i>	<i>IV</i>
		<i>Sample by Household</i>			
		<i>Non-Collat. Debt Level</i>			
Variable		All	No Debt	Positive Debt	<i>p</i> -value of difference
Age of the Household Head		31.6	32.3	31.1	<0.01
<i>Percent of households with...</i>					
Sex of Head = Male		0.83	0.84	0.83	0.58
Race of Head = Black		0.12	0.16	0.09	<0.01
Married		0.55	0.50	0.58	<0.01
Non-Collateralized Debt		0.62	–	–	–
Head’s 1998 Labor Income	Mean	33,600	35,400	32,500	<0.05
	Median	27,400	27,400	27,800	0.40
<i>Numbers below all divided by income</i>					
Head’s Total SS Contributions	Mean	0.85	0.91	0.81	<0.01
	Median	0.64	0.68	0.62	0.02
Household Non-Coll. Debt	Mean	0.36	–	0.58	–
	Median	0.05	–	0.22	–
Household Total Net-Worth	Mean	2.02	3.07	1.80	<0.01
	Median	0.66	1.00	0.48	<0.01
Household Net Liquid Assets	Mean	0.24	0.78	-0.08	<0.01
	Median	0	0.08	-0.08	<0.01
Household Net Financial Assets	Mean	0.04	0.50	-0.24	<0.01
	Median	0	0	-0.16	<0.01
Sample Size		2,077	844	1,233	

Table 2: Portfolio Composition of PSID Households under 40, by Debt Holding. See notes to Table 1. [tabk2]

	<i>I</i>	<i>II</i>	<i>III</i>	<i>IV</i>
	<i>Sample by Household</i>			
	<i>Debt Level</i>			
Variable	All	No Debt	Positive Debt	<i>p</i> -value of difference
Owning a Home	0.44	0.44	0.44	0.81
Having Positive Cash	0.83	0.74	0.89	<0.01
Owning Any Stocks	0.22	0.22	0.22	0.99
Owning Any Bonds	0.17	0.15	0.18	0.06
Owning Cash, Stocks or Bonds	0.85	0.77	0.90	<0.01
Owning Stocks or Bonds	0.32	0.30	0.33	0.26
Owing a Business	0.09	0.10	0.09	0.91
with Negative Net Worth	0.16	0.01	0.26	<0.01
with Negative Net Liquid Assets	0.39	0	0.63	<0.01
with Negative Net Financial Assets	0.50	0	0.82	<0.01
Sample Size	2,077	844	1,233	

Table 3: Descriptive Statistics for Single PSID Households under 40, by Debt Holding See Notes to Table 1. [tabk3]

Variable	<i>Sample by Household Debt Level</i>				<i>p</i> -value of difference
	<i>I</i>	<i>II</i>	<i>III</i>	<i>IV</i>	
	All	No Debt	Positive Debt		
Age of the Household Head	30.3	31.2	29.7		<0.01
<i>Percent with</i>					
Sex of Head = Male	0.63	0.68	0.60		0.02
Race of Head = Black	0.18	0.24	0.14		<0.01
Non-Collateralized Debt	0.58	–	–		–
Head's 1998 Labor Income	Mean	27,000	28,300	26,300	0.21
	Median	24,100	23,600	24,500	0.42
<i>Numbers below all divided by income</i>					
Head's Total SS Contributions	Mean	0.81	0.87	0.77	0.09
	Median	0.61	0.66	0.59	0.09
Household Non-Coll. Debt	Mean	0.46	–	0.79	–
	Median	0.04	–	0.25	–
Household Total Net-Worth	Mean	1.04	1.70	0.56	<0.01
	Median	0.31	0.49	0.15	<0.01
Household Net Liquid Assets	Mean	0.00	0.58	-0.42	<0.01
	Median	0	0.05	-0.14	<0.01
Household Net Financial Assets	Mean	-0.18	0.32	-0.54	<0.01
	Median	-0.01	0	-0.20	<0.01
Sample Size	850	379	471		

Table 4: Portfolio Composition of Single PSID Households under 40, by Debt Holding. See Notes to Table 1. [tabk4]

	<i>I</i>	<i>II</i>	<i>III</i>	<i>IV</i>
	<i>Sample by Household</i>			
	<i>Debt Level</i>			
<i>Percent...</i>	All	No Debt	Positive Debt	<i>p</i> -value of difference
Owning a Home	0.22	0.24	0.21	0.44
Having Positive Cash	0.77	0.65	0.86	<0.01
Owning Any Stocks	0.17	0.18	0.17	0.88
Owning Any Bonds	0.13	0.13	0.13	0.97
Owning Cash, Stocks or Bonds	0.79	0.69	0.87	<0.01
Owning Stocks or Bonds	0.26	0.26	0.26	0.99
Owing a Business	0.06	0.06	0.05	0.62
with Negative Net Worth	0.22	0.01	0.37	<0.01
with Negative Net Liquid Assets	0.40	0	0.69	<0.01
with Negative Net Financial Assets	0.50	0	0.88	<0.01
Sample Size	850	379	471	

Table 5: Median Ratio of Net Financial Assets to Income, by Age Range. See Notes to Table 1. [tabk5]

Age Range	Debt/Income by %ile			SS Wealth /Income
	50	75	90	
22 – 24	0.13	0.45	1.54	0.22
25 – 27	0.14	0.54	1.11	0.33
28 – 30	0.12	0.37	0.90	0.60
31 – 33	0.03	0.27	0.70	0.70
34 – 36	0.02	0.20	0.54	1.07
37 – 39	0.03	0.20	0.56	1.20

Table 6: Distribution of Non-Collateralized Debt to Income Before and After Adjusting for Social Security Wealth Notes – Data is from the 1999 PSID. Sample in columns I and II includes 1) all unmarried households with heads between the age of 22 and 39 and 2) married households with heads between the age of 22 and 39 where the head also has a complete earnings history (2,077 households). Sample in column III includes all unmarried households with heads between the age of 22 and 39 (850 households). Social Security augmented debt is defined as non-collateralized debt less total accumulated Social Security contributions of the head. Social Security augmented debt is constructed to be non-negative. In other words, if Social Security contributions exceed the household’s non-collateralized debt, Social Security augmented debt is set equal to zero. [tabk6]

	<i>I</i>	<i>II</i>	<i>III</i>
	All Households	All Households With Debt	Single Households With Debt
<u>Panel A</u>			
<i>Percent With Positive Non-Collateralized Debt</i>			
	0.62	1.00	1.00
<i>Distribution of Debt to Income Ratio</i>			
25 th Percentile	0.00	0.08	0.09
50 th Percentile	0.05	0.22	0.25
75 th Percentile	0.29	0.53	0.60
90 th Percentile	0.84	1.20	1.72
Mean	0.36	0.58	0.79
<u>Panel B</u>			
<i>Percent with Positive Social Security Augmented Debt</i>			
	0.17	0.28	0.31
<i>Distribution of Soc Security Augmented Debt to Income Ratio</i>			
25 th Percentile	0	0	0
50 th Percentile	0	0	0
75 th Percentile	0	0.04	0.14
90 th Percentile	0.33	0.74	0.97
Mean	0.19	0.32	0.50

Table 7: Parameter Settings

Parameter	Baseline	Alternative values
Relative Risk Aversion	3	
Annual Discount factor	0.97	0.98
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 8: Household consumption and portfolio choice over the life-cycle as predicted by the model under various assumptions. Equity benefit measures the percentage increase in certain equivalent consumption yielded by investment in equity. [tab1.11] [ss.11.txt]

#	β	Profile	Equity benefit	22-24	25-27	28-30	31-33	34-36	37-39	22-39	
1	0.97	Baseline	2.39	Consumption	16.8	18.0	19.2	20.5	21.8	23.3	19.9
				Stocks+Bonds/Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00
				Debt/Y	0.27	0.40	0.45	0.42	0.32	0.16	0.33
				SS Wealth/Y	0.30	0.58	0.85	1.12	1.40	1.67	1.06
2	0.97	High School	2.55	Consumption	16.4	17.5	18.7	20.0	21.3	22.7	19.4
				Stocks+Bonds/Y	0.00	0.00	0.00	0.00	0.00	0.02	0.00
				Debt/Y	0.04	0.09	0.13	0.14	0.10	0.01	0.08
				SS Wealth/Y	0.30	0.60	0.90	1.18	1.45	1.72	1.09
3	0.97	College	3.03	Consumption	19.9	21.2	22.6	24.1	25.7	27.0	23.4
				Stocks+Bonds/Y	0.00	0.00	0.00	0.00	0.02	0.17	0.04
				Debt/Y	0.50	0.48	0.33	0.15	0.01	0.00	0.22
				SS Wealth/Y	0.28	0.53	0.81	1.12	1.45	1.77	1.06
4	0.97	Flat	8.55	Consumption	22.8	23.8	24.7	25.7	26.8	27.9	25.3
				Stocks+Bonds/Y	0.46	0.94	1.41	1.90	2.42	2.96	1.68
				Debt/Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00
				SS Wealth/Y	0.31	0.64	0.99	1.36	1.75	2.16	1.20
5	0.98	Baseline	3.62	Consumption	16.4	17.7	19.0	20.5	22.1	23.4	19.8
				Stocks+Bonds/Y	0.00	0.00	0.00	0.00	0.03	0.23	0.05
				Debt/Y	0.17	0.23	0.21	0.12	0.01	0.00	0.11
				SS Wealth/Y	0.30	0.58	0.85	1.12	1.40	1.67	1.06

Table 9: The effects of social security on consumption and portfolio choice over the life cycle as predicted by our model. Scenario is baseline (see Table 7). [tab1.9] [ss.9.txt]

#	Init. Wealth/Income	Utility	22	26	30	34	38	66
	Non-SS	SS						
1			16.5	18.0	19.6	21.4	23.3	34.6
			0.0	0.0	0.0	0.0	0.0	362.0
			3.3	7.9	10.1	9.2	4.6	0.0
		0.0	(0.3)	(0.8)	(1.0)	(0.9)	(0.5)	(0.0)
			3.3	11.2	21.0	32.8	46.7	0.0
2			(0.1)	(0.2)	(0.4)	(0.7)	(0.9)	(0.0)
			18.0	19.7	21.5	23.4	25.2	38.5
			0.0	0.0	0.0	0.0	7.5	401.3
			3.2	6.9	7.1	2.9	0.0	0.0
		0.0	(0.3)	(0.7)	(0.7)	(0.3)	(0.0)	(0.0)
3			0.0	0.0	0.0	0.0	0.0	0.0
			(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
			-	-	19.7	21.5	23.5	34.4
			-	-	0.0	0.0	0.0	360.3
			-	-	8.5	7.6	3.0	0.0
4			-	-	(0.8)	(0.8)	(0.3)	(0.0)
			-	-	15.8	27.1	40.5	0.0
			-	-	(0.3)	(0.5)	(0.8)	(0.0)
			-	-	23.0	24.5	26.0	39.1
			-	-	4.5	7.6	18.1	408.4
5			-	-	0.0	0.0	0.0	0.0
			-	-	(0.0)	(0.0)	(0.0)	(0.0)
			-	-	0.0	0.0	0.0	0.0
		0.0	-	-	(0.0)	(0.0)	(0.0)	(0.0)
			-	-	(0.0)	(0.0)	(0.0)	(0.0)

Table 10: The effects of allowing households to use social security wealth to pay off debt. Withholding is the percentage of income withheld for social security contributions. ILR is the internal lending rate (See Section 3.1 for an explanation.) Private saving equals stocks plus bonds minus debt. Scenario is baseline (see Table 7). Replacement 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 certain equivalent consumption (see equation (1)). Portfolio loss is the percentage reduction in utility of the scenario as compared to the no social security benchmark (panel 6 in this case). [tab1.7] [ss.7.txt]

#	Age 30 Wealth		Parameters of SS System			Age:
	Non-SS	SS	With- holding	ILR	Repl. Rate	
1	-0.4	0.6	10.0	2	42	Consumption 19.7 22.5 26.8 30.5 34.4 Private saving -8.5 -5.8 23.4 95.6 360.3
2	0.0	0.2	11.7	2	43	Consumption 20.3 23.1 26.9 30.6 34.4 Private saving -0.4 0.6 27.8 96.0 360.5
3	-0.4	0.6	10.0	3	51	Consumption 19.8 22.6 27.1 31.0 34.9 Private saving -8.5 -6.7 18.3 79.5 364.9
4	-0.4	0.6	10.0	4	64	Consumption 19.9 22.7 27.5 31.6 35.6 Private saving -8.7 -7.9 12.0 59.8 372.7
5	-0.4	0.6	10.0	5	80	Consumption 20.1 22.8 28.0 32.3 36.9 Private saving -8.8 -9.2 5.2 36.7 386.3
6	0.2	0.0	0.0	2	0	Consumption 23.0 25.2 29.0 32.8 39.1 Private saving 4.5 11.8 71.5 214.3 408.4

Table 11: Effects of exemptions. Scenario is baseline (see Table 7). See Table 10 for definitions. [tab1.8] [ss.8.txt]

#	Parameters of SS System				Age:	26	31	36	51	66
	Start Age	With-holding	ILR	Repl. Rate						
1	21	10.0	2	43	Consumption	18.0	20.0	22.3	28.6	34.6
					Private saving	-7.9	-10.2	-7.4	51.2	362.0
2	30	12.3	2	43	Consumption	19.0	20.9	23.0	28.5	34.4
					Private saving	-2.0	0.2	0.8	56.2	360.1
3	40	17.3	2	43	Consumption	19.4	21.6	23.6	28.8	34.7
					Private saving	-5.0	-1.6	12.2	75.2	363.6
4	50	29.8	2	43	Consumption	19.6	21.8	24.1	29.3	35.5
					Private saving	-6.1	-4.5	5.5	121.7	372.4
5	21	10.0	3	54	Consumption	18.0	20.1	22.4	29.0	35.0
					Private saving	-8.3	-11.1	-9.3	39.9	366.4
6	21	10.0	4	68	Consumption	18.1	20.1	22.4	29.6	35.9
					Private saving	-8.7	-12.0	-11.2	27.2	375.8
7	21	10.0	5	88	Consumption	18.1	20.2	22.5	30.3	37.5
					Private saving	-9.2	-13.2	-13.5	12.1	392.6
8	21	8.0	3	43	Consumption	18.4	20.4	22.8	29.3	35.4
					Private saving	-8.1	-10.2	-7.3	53.5	369.9
9	21	6.3	4	43	Consumption	18.7	20.8	23.2	29.8	36.0
					Private saving	-8.2	-10.4	-7.3	55.4	376.5
10	21	4.8	5	42	Consumption	19.0	21.2	23.6	30.1	36.5
					Private saving	-8.4	-10.6	-7.6	57.5	381.4
11	21	0.0	2	0	Consumption	19.7	21.9	24.4	30.0	38.5
					Private saving	-6.9	-6.5	1.3	117.9	401.3

Table 12: Effects of exemptions. Scenario is baseline (see Table 7) except that households cannot trade equity. See Table 10 for definitions. [tab1.12]
[ss.12.txt]

#	Parameters of SS System				Age:	26	31	36	51	66
	Start Age	With-holding	ILR	Repl. Rate						
1	21	10.0	2	43	Consumption	18.1	20.2	22.5	27.3	25.9
					Private saving	-9.0	-12.8	-12.8	23.4	307.8
2	30	12.3	2	43	Consumption	19.1	20.9	23.3	26.9	25.5
					Private saving	-2.8	-1.3	-2.7	24.1	302.6
3	40	17.3	2	43	Consumption	19.9	22.1	24.6	25.6	24.3
					Private saving	-8.1	-9.4	-4.5	20.2	288.7
4	21	0.0	2	0	Consumption	19.9	22.2	24.7	25.6	24.2
					Private saving	-8.4	-10.3	-6.3	92.4	287.4

Table 13: Effects of changes in the Internal Borrowing Rate (IBR) (See Section 3.1 for an explanation) for an Age-40 exemption. Scenario is baseline (see Table 7) except where noted. See Table 10 for definitions.[tab1.10] [ss.10.txt]

#	Parameters of SS System				Age:	26	31	36	51	66
	Start Age	With-holding	IBR	Repl. Rate						
1	21	12.0	2	43	Consumption	18.0	20.0	22.3	28.6	34.6
					Private saving	-7.9	-10.2	-7.4	51.2	362.0
2	40	19.3	2	43	Consumption	19.4	21.6	23.6	28.8	34.7
					Private saving	-5.0	-1.6	12.2	75.2	363.6
3	40	23.0	4	43	Consumption	19.3	21.5	23.2	28.2	34.0
					Private saving	-4.2	0.4	17.1	79.1	356.7
4	40	25.4	5	43	Consumption	19.2	21.3	23.0	27.9	33.6
					Private saving	-3.6	1.9	20.5	81.8	352.4
5	40	28.3	6	43	Consumption	19.1	21.1	22.7	27.5	33.2
					Private saving	-2.9	3.9	24.8	85.5	347.6
6	40	35.8	8	43	Consumption	18.9	20.6	22.1	26.6	32.1
					Private saving	-1.0	9.2	36.3	95.3	332.6
7	21	2.0	2	0	Consumption	19.7	21.9	24.4	30.0	38.5
					Private saving	-6.9	-6.5	1.3	117.9	401.3

Figure 1: Expected labor income profiles. figx7

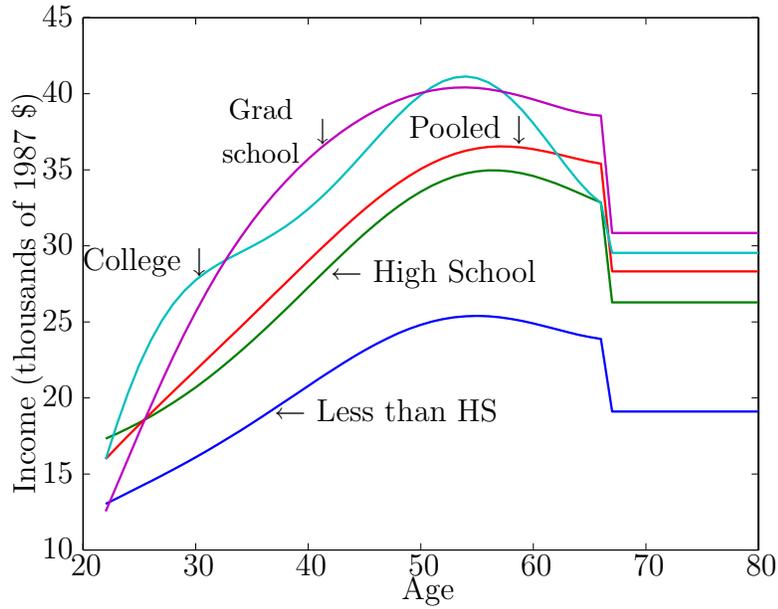


Figure 2: Excess saving over the life cycle for various different social security schemes. Excess saving equals the difference between total saving (including social security) under the listed scheme and total saving under the optimal scenario for the household in which there is no social security. Age 30 refers to an exemption ending at age 30 and so on. Existing is the existing system. [figa1]

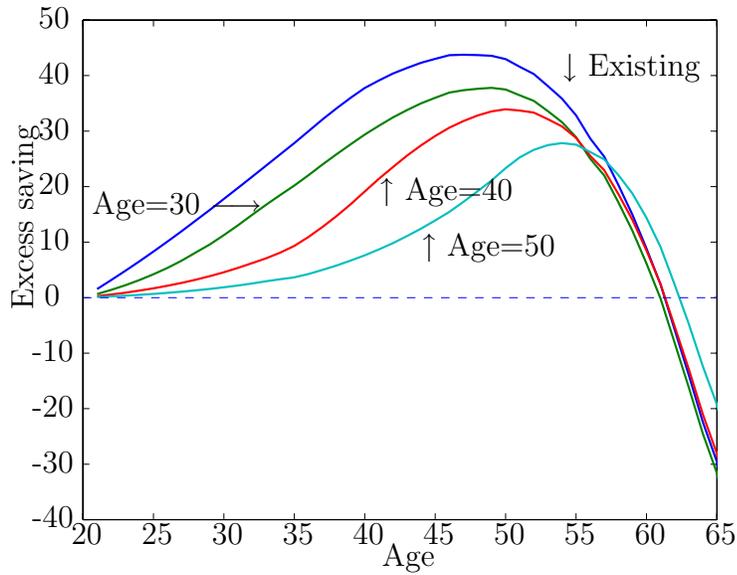


Figure 3: Portfolio allocation to equity under various different social security schemes. “Optimal” refers to the optimal scenario for the household in which there is no social security. For other definitions, see Figure 2. [figa2]

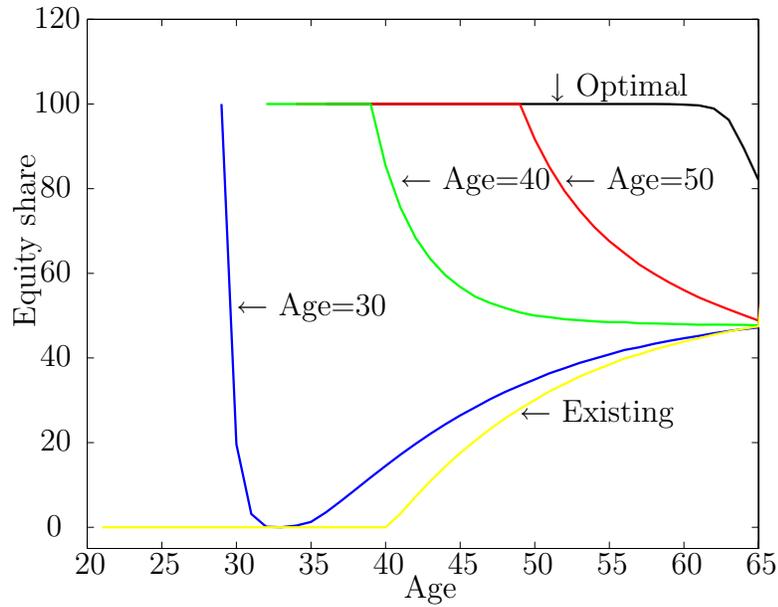


Figure 4: Excess private saving over the life cycle. Excess *private* saving equals the difference between private saving (total saving minus social security) under the listed scheme and total saving under the optimal scenario for the household in which there is no social security. [figa3]

