Public Policy
Discussion Papers
Federal Reserve Bank of Boston

No. 04-10

Social Security and Unsecured Debt

Erik Hurst and Paul Willen

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 consider a life-cycle model with "tempted" households, who find it impossible to commit to an optimal consumption plan and "disciplined" households who have no such problem, and we explore ways to reduce this inefficiency. We show that allowing households to use social security wealth to pay off debt while exempting young households from social security contributions (but in both cases requiring higher contributions later) leads to increases in welfare for both types of households and, for disciplined households, to significant increases in consumption and saving and reductions in debt.

JEL Classifications: D91, G11, H31

Erik Hurst is an Associate Professor of Economics at the University of Chicago's Graduate School of Business and a research fellow at the NBER. Paul Willen is a Senior Economist at the Federal Reserve Bank of Boston and a research fellow at the NBER. Their e-mail addresses are: erik.hurst@gsb.uchicago.edu and paul.willen@bos.frb.org, respectively.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at http://www.bos.frb.org/economic/ppdp/index.htm.

The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Boston or the Federal Reserve System.

We thank Charlie Brown, Jeff Liebman, Borja Larrain, Olivia Mitchell, Geoff Tootell, Kent Smetters, Steve Zeldes, 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. Both authors gratefully acknowledge research support from the Graduate School of Business at the University of Chicago.

This version: December 2004

1 Introduction

The starting point for our analysis is the observation that households currently have about $700 billion in unsecured debt on which they pay roughly 10 percent interest and $11 trillion of social security wealth on which they earn less than 2 percent interest.[1] As a nation, we are apparently borrowing on credit cards and saving in a passbook savings account, something that Gross and Souleles (2002) describe as "puzzling, apparently inconsistent with no-arbitrage and thus inconsistent with any conventional model."

We focus on the old-age portion of social security and explore this topic in three steps. First, in Section 2, we examine the distribution across individuals of non-collateralized debt and social security wealth. In an economy with heterogenous 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 (PSID) to show that this is generally not the case. In our sample of households under the age of 40, 62 percent have unsecured debt.[2] We show that if households could access their social security wealth to pay off debt, only 17 percent would still have 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.

In Section 3, we construct a dynamic life-cycle portfolio choice model. We build on utility theory developed by Gul and Pesendorfer (2004a) and assume that the world is populated by two types of households, differing only in their attitude to current consumption: *tempted* households, which care about the future but face an overwhelming desire to consume all available resources in the current period, and *disciplined* households, which face no such desire. Such preferences are consistent with the axiomatic basis established in Gul and Pesendorfer's paper. Tempted households suffer from many of the behavioral inconsistencies documented by Laibson (1997) and share the preference for commitment of households that use the quasi-hyperbolic discounting model developed by Laibson and simulated in Laibson, Repetto, and

[1] Our measure of unsecured debt is "consumer revolving credit" from the Federal Reserve Board – which was $725.0 billion in May 2003. Davis, Kubler, and Willen (2003) argue that the interest rate on unsecured debt, taking default into account, is roughly 10 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 Ranguelova (2001). Leimer (1994) calculates that the internal rate of return on social security contributions is 1.7 percent.

[2] In an abuse of language, we refer to a household headed by a 40-year old as "a 40-year-old household."

Tobacman (1998).[3] We show that tempted households consume their income each period and thus follow the "rule-of-thumb" of Campbell and Mankiw (1989). For the disciplined households, we adapt a model developed by Davis, Kubler, and Willen (2003). 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 is, the lending rate). Such an assumption is consistent with the pattern of observed interest rates.[4] We show that our parameterization of this model can roughly match the life-cycle borrowing behavior documented in Section 2. Evidence from the consumption literature suggests that households roughly break into these two categories, especially with respect to retirement.[5]

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 of Hubbard and Judd (1987) and exempt young households from social security contributions. Under both proposals, households 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 certain equivalent consumption.[6] Table 1 summarizes our key results. For example, moving from the existing system to one in which households under 30 are exempt from contributing to social security (but are forced to make up the taxes later in life) raises certain-equivalent consumption by 3.4 percent for disciplined households and 3.3 percent for tempted households.

Our analysis also yields valuable insights for the current debate on reforming social security. We argue that, as a forced saving scheme, social security suffers from two serious flaws: It forces people to save at points in the life cycle when even disciplined households don't want to; and it forces everyone to invest everything in low-yielding assets. It's easy to see how our exemptions address the former problem, but we show that exemptions address the latter problem, as well. Specifically, we show that, with the exemptions, disciplined households come closer to approximating

[3] See Gul and Pesendorfer (2004b) for a "low-brow" discussion of their results which shows how their preferences can generate the preference reversals usually used to motivate hyperbolic discounting.

[4] See Davis et al., 2003.

[5] See the discussion in Section 1.1 below.

[6] Certain-equivalent consumption is the level of constant consumption over the lifecycle which yields the same utility as the risky stream the household actually receives.

the optimal allocation across risky and riskless assets without actually investing any of their social security wealth in risky assets, a fact we illustrate in Section 4. Thus, exemptions allow us to address one 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 effectively lends money to households. We call the interest rate on such lends 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. Table 1 shows that a combination of the 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 types of households. We also show that raising the IBR would significantly improve social security finances. In other words, we show that the government could borrow at 2 percent, lend at 5 percent, and still make households better off!

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

First, we make no assumption 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 estimate 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.[7] 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

[7]See Abel (1986) and Eckstein, Eichenbaum, and Peled (1985), among others.

eliminating social security altogether – has any effect on the retirement portion of social security. Thus, the benefits shown in Table 1 and elsewhere in the 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 (2002), 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, 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. So it is particularly ironic that the contribution structure does precisely the opposite at certain points in the life cycle. Second, a progressive income tax approximates the ideal life-cycle structure by lowering tax rates when income is low. Since social security taxation is, in fact, regressive, not progressive, 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. 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 only requires a marginal-tax-rate increase 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.

Second, we have constructed a partial-equilibrium model. Obviously, if a policy change has significant partial-equilibrium effects, one imagines 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. (See Bohn 1999 and Diamond and Geanakoplos 2002 for examples.)

Before proceeding with the main body of the paper, we conduct a brief literature review.

1.1 Literature review

This paper builds on earlier work in many fields, spanning consumption and portfolio choice and public finance.

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 three particularly relevant papers. Feldstein and Ranguelova (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 in two key respects. First, we analyze the effects of exemptions on households that don't plan for the future, our "tempted" households. And second, households in our model can invest in stock and borrow – opportunities not offered to households in Hubbard and Judd's model, opportunities that, one would expect, worked make forced investment in social security both more costly and easier to circumvent. The finding that exemptions still offer such large benefits is therefore quite surprising.

In a series of papers, Laibson argues that the standard representation of utility with exponential discounting cannot account for key pieces of survey evidence or for the apparent desire of households to commit to future consumption behavior – an issue of particular relevance to social security, which Akerlof (1998), among others, justified precisely because of its commitment aspects. Laibson argues that the

quasi-hyperbolic (or quasi-geometric, according to Krusell and Smith 2003) alternative generates behavior in line with survey and other evidence. Gul and Pesendorfer (2001,2004a) propose an alternative model that could rationalize these facts but which one can solve using standard dynamic programming techniques, and we take advantage of their insights here.[8]

Two notable papers address the issue of commitment in general and its relationship with social security in particular in the context of models with hyperbolic discounting. Laibson, Repetto, and Tobacman (1998) show that one commitment device, 401 (k) plans, raises utility much more for hyperbolic than for exponential households in a partial-equilibrium model. Our findings for "tempted" households are completely consistent with the findings of Laibson et al. Imrohoroglu, Imrohoroglu, and Joines (2003) also look at quasi-hyperbolic preferences and find that social security is a reasonable substitute for perfect commitment in partial equilibrium but a poor substitute in general equilibrium, where depressed saving leads to a much lower capital stock and thus to lower income in equilibrium.

Gross 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.

Researchers have recently started to focus attention on the effects of borrowing limitations on life-cycle portfolio choice.[9] We build on this literature by extending a model of Davis, Kubler, and Willen (2003). They 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 both that a wedge is empirically evident and that a model incorporating a wedge leads to a more realistic profile of life-cycle portfolio holdings. 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.

Altig and Davis (1992) consider the effects of inserting a wedge between borrowing and lending rates in an overlapping generations model. They find that the wedge dramatically affects the timing of bequests. Our finding, that a wedge affects the optimal timing of contributions, is similar in spirit.

[8]For a comparison of the hyperbolic and Gul-Pesendorfer models, see Gul and Pesendorfer (2004b).

[9]Some 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.

Finally, recent work in the consumption literature motivates our decision to categorize consumers as either disciplined or tempted. Researchers have shown that consumers whose behavior roughly matches that implied by the two definitions 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 a disciplined model describes consumption behavior at retirement for roughly 80 percent of the population, while what we call a tempted model describes the behavior of the remaining 20 percent.

2 Empirical facts about borrowing and social security

2.1 Data

The *PSID* is a large, nationally representative survey, started in 1966, that tracks social and economic variables of a given household over time. Each year, the survey gathers 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 reported 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 reported the value of their main home, the value of their outstanding mortgage debt, 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 unsecured debt (including store and credit card debt, student loans, and other personal loans).

The time-series aspect of the PSID makes it ideal for measuring social security wealth. While the PSID does not have actual social security records for each respondent, social security wealth can be computed using the household's detailed earnings history coupled with social security tax tables.[10]

[10]See `http://www.ssa.gov/OACT/`.

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 approach. 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 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 40. We focus on younger households since most unsecured borrowing occurs among younger households. Unfortunately, we do not have full earnings histories for all household heads in the sample. To see why, note two things. First, 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. Second, given PSID definitions, the male is classified as the household head for all married and cohabiting couples. As a result, for some male PSID heads, earnings histories extend back only to the date that they married into the survey. Rather than imputing the missing income history for these household heads, we restrict 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 descendents who are equally likely to be men or women, our restriction does not bias our analysis in any way. In summary, our 1999 PSID sample includes all unmarried households with heads between the ages of 22 and 40 and married households with heads between the ages of 22 and 40 where the head also has a complete earnings history. The resulting sample size is 2,077 households. For robustness, we also examine only single-headed households where the head is between 22 and 40. This latter sample includeds 850 households.

2.2 Discussion of data

Our analysis of the data reveals three key facts. First, young households have large amounts of unsecured debt. Table 2 shows some descriptive statistics about our sample. 62 percent of households under the age of 40 had positive amounts of unsecured debt. The median and mean amount of debt held for households in our sample were, respectively, $1,200 and $6,400. For those that held positive amounts of debt, the

median debt held was \$5,600, while the mean was \$10,500.[11] The results are equally striking when comparing total household debt in 1999 with the household head's current 1999 income. The median debt-to-income ratio was 5 percent; the mean was 36 percent.[12] Households that had positive unsecured debt tended to be slightly younger, were 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, unsecured debt. Low-income households are unable to obtain debt, despite their desire to borrow, while high-income households have less need to borrow. Among households that had some debt – that is, 62 percent of our sample – the median debt-to-income ratio was 22 percent and the mean was 58 percent.

If individuals are constrained by social security contributions, we would expect that households with debt would 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 Treasuries less unsecured debt.[13] Second, we create a measure called "net liquid assets" by subtracting illiquid assets from our measure of net worth. We define illiquid assets as business equity, housing equity, and vehicle equity. Finally, realizing that cash, saving accounts, and checking accounts may be used to make monthly purchases, we subtract these resources from net liquid assets to get a third measure which we call "net financial wealth." 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. Table 3 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

[11] All dollar amounts in this paper are in 1996 constant dollars, unless otherwise noted.

[12] Unlike 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 anytime 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, who carry balances forward from month to month). Of those households who 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 balance on a monthly basis (Gross and Souleles, 2002).

[13] The 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 to the SCF wealth data (Juster, Smith, and Stafford, 1999).

vehicle equity, for example). If we look at "net financial assets," we get a completely different picture. Table 3 shows that 50 percent of the households in the sample have negative net financial assets, and 82 percent of the households with unsecured debt have negative net financial assets. Our sample does indicate that some households do have positive holdings of stocks and bonds at the same time that they have unsecured debt. To see this, look at Table 2, which shows that the median household with positive debt has net financial assets equal to minus 16 percent of income but debt equal to 22 percent of income. In other words, if households with positive debt used their financial assets to pay off their debt, they could reduce their debt load from 22 percent of income to 16 percent. However, the PSID data tell us neither the interest rate on the debt nor the return on the assets. If the return on stocks is sufficiently high or the interest on the debt is sufficiently low, such debt reduction does not make sense. By contrast, the interest rate on unsecured debt always exceeds the interest rate on social security contributions.

The life-cycle profile of debt is as expected. Young households have steep income profiles, on average. As a result, the Permanent Income Hypothesis would predict that young households should be indebted. As they age, and their income profiles flatten out, the propensity to be in debt should diminish. These predictions are borne out in the data. Table 4 shows that the median debt-to-income ratio rises slightly from 13 percent between the ages of 22 and 24, to 14 percent between the ages of 25 and 27, before falling sharply to 3 percent between the ages of 37 and 39. Variations for the 75th percentile household are more dramatic. Table 4 shows that the debt-to-income ratio for the 75th-percentile household rises from 45 percent for household heads aged between 22 and 24 and to 54 percent for households aged 25 to 27, before falling to 20 percent for households aged 37 to 39.

Indebtedness is highly skewed. In the whole sample, the median household debt-to-income ratio is just 5 percent. Table 5 shows that a quarter of households 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 distribution is still quite skewed. Twenty-five percent of households have debt-to-income ratios of less than 8 percent. However, 25 percent also have debt-to-income ratios of more than 53 percent. And 10 percent have debt-to-income ratios that exceed 120 percent.

Social security wealth is considerable for young households. Table 2 shows that total social security contributions for the with-debt and the without-debt samples are broadly similar. For the median household in our sample with debt, the social security wealth-to-income ratio is 62 percent, and for the mean household, it is 81

percent. As expected, given social security rules and life-cycle income profiles, social security wealth rises rapidly over the life cycle. Table 4 shows that the social security wealth-to-income ratio grows from 22 percent of income for household heads aged between 22 and 24 to 120 percent of income for households aged 37 to 39.

Most interestingly, the data bear out our basic claim that households could wipe out much of their unsecured debt if they had access to their social security wealth. To measure the effects of social security wealth on debt, we construct "social security augmented debt" by subtracting social security wealth from household debt. If social security wealth exceeds debt for a particular household, then social security augmented debt equals zero. Table 5 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 that have debt, 72 percent can eliminate it when they have access to social security wealth. For those that cannot eliminate their debt, the level of indebtedness falls dramatically. The 90th-percentile household has debt equal to 84 percent of income; after gaining access to social security wealth, the 90th-percentile household has debt equal to 33 percent of income. 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. It should be noted that these numbers should be seen as upper bounds. Given the structure of the PSID, we can measure the actual social security wealth for only one member of the household. Debt, however, is reported for all household members.[14]

In terms of aggregates, there is a large decline in total debt and total interest paid for the U.S. economy when households under 40 are allowed to use social security wealth to pay off non-collateralized debt. Table 6 shows some aggregate calculations. Using PSID sample weights and census data, we calculate that allowing households to pay off non-collateralized debt with social security wealth would reduce total household debt by more than $100 billion and reduce total annual household interest payments by $11.5 billion per year – or about $500 for every household that has non-collateralized debt.

[14]We 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 the spouse's social security wealth falls to 14 percent.

3 A life-cycle model

We now construct a model of life-cycle consumption with social security. 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.[15] We assume that labor is supplied inelastically and that household income evolves deterministically over the life cycle.

Households can trade three financial assets: They can buy equity with stochastic net return $\tilde{r}_E$ and bonds at a net risk-free rate r_L, and they can 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.

We model household utility using a version of Gul and Pesendorfer's (2004a) dynamic self-control model. It is often argued that social security exists because many people lack the discipline needed to save for retirement. Gul and Pesendorfer preferences allow households to exhibit such problems but – unlike hyperbolic discounting models and other models that also allow for discipline problems – their model can be solved using standard dynamic programming techniques and, more importantly for our purpose, allow for straightforward welfare comparisons.[16]

Let x equal current financial wealth, y equal current income, and c equal current consumption; and let primes denote one-period-ahead values. Following Gul and Pesendorfer, we set household utility recursively:

$$W(x, y) = \max_{c \in C} \{u(c) + \delta W(x', y') + \phi u(c)\} - \phi u(x + y), \tag{1}$$

where u is an isoelastic period utility function with relative risk-aversion coefficient γ. We examine two types. We call households with $\phi = 0$ "disciplined" households, and it's easy to see that with $\phi = 0$, equation (1) is a standard von-Neumann-Morgenstern utility function. At the other extreme, we define "tempted" households, with $\gamma = \infty$, which we show below implies "rule-of-thumb" behavior.

We solve the model computationally using methods developed in Judd, Kubler, and Schmedders (2002). For details on our numerical solution, see the appendix to Davis, Kubler, and Willen (2003). Table 7 lists key features of our parameterization of the model. We draw the reader's attention to several aspects of our parameterization.

[15] For our results on tempted households, we actually need to assume that households have an arbitrarily small but positive level of income in retirement.

[16] See see Akerlof (1998) and Laibson, Repetto, and Tobacman (1998) for justifications of forced savings schemes using time-inconsistent agents.

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, the assumption of inelastic labor supply is a strong one, and we revisit the issue in the last section. Our solution for disciplined households allows for stochastic variation in labor income, but we go with a deterministic specification for two reasons. First, our focus is on the life-cycle aspects of borrowing and saving, not the high-frequency variation. Second, Davis, Kubler, and Willen (2003) show that the predictions of our model for wealth accumulation for relatively impatient households with labor-income risk are similar to those of our model with more patient households with no labor-income risk. To account for this, we estimate the model with different time discount rates.

Third, Davis, Kubler, and Willen (2003) present evidence that the borrowing rate exceeds the riskless lending rate by 10 percentage points. Of those 10 percentage points, charge-offs for uncollected loans (that is, defaults) 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, which yields a borrowing rate of 10 percent. Our model is partial equilibrium, so we do not attempt to explain the origins of the wedge between borrowing and lending rates. However, Dubey, Geanakoplos, and Shubik (2003) 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.

Fourth, 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 (2003). 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.[17]

3.1 Social security system

Our social security system works in the following way. Households pay a proportional tax while working. Upon retirement, disciplined households receive a lump-sum pay-

[17]Gokhale, 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), which 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 prior to age 30, these effects will be small.

ment equal to the value of an annuity paying a fixed fraction of the average of the highest 35 years' of income. Tempted households receive the annuity.

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. Thus, household contributions to social security are a declining fraction of income. In our model, there is no cap, but this omission actually strengthens our results. We argue below that 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 *raising* the tax rate as income goes down.

How do benefits differ between our social security system and the real one? 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 one's income that one receives in retirement depends on the level of one's 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. In essence, the social security system provides 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. However, in all our policy experiments (aside from the straw man of eliminating social security altogether), the benefits portion of the social security system remains unchanged.

Second, we give a lump sum, rather than an annuity, to disciplined 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. As noted above, 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 what-

ever welfare benefits an annuity confers are preserved in our policy experiments. 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 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 are identical under all plans. Below, we discuss this adjustment in depth.

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 (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), 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 percent We round up and assume that the internal lending rate is 2.0 percent. In our baseline scenario, this implies that the replacement rate in retirement equals around 43 percent.[18]

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 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 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: We allow households 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 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

[18]The Social Security Administration in 2003 states that their replacement rate has been stable at about 42 percent for the last 20 years (See http://www.ssa.gov/kc/fact_sheet_14_exp.htm).

the internal lending rate. The existing social security system does allow some of this sort of "borrowing": 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.[19]

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 disciplined households, a forced saving scheme as defined above will either have no effect on welfare or make households worse off. If disciplined 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 disciplined households face restrictions on borrowing, social security will typically make them worse off, because they can no longer undo it. Thus, for disciplined households, borrowing mitigates the costs of a forced saving scheme.

For tempted households, if the household can borrow and lend unlimited amounts at the riskless rate, then it 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 Basic results of the disciplined model

In this section, we outline the basic implications for life-cycle consumption and portfolio choice of our model for disciplined households. For a thorough discussion of the results of the model, see Davis, Kubler, and Willen (2003). In the discussion that follows we 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 for disciplined households all wealth in retirement is "private saving." Four aspects of the solution are relevant to our discussion of social security.

[19] Authors' calculation using benefit and life expectancy information from the Social Security Administration web site.

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. 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. Conversely, suppose we raise the borrowing rate: Equity holdings will also go up. 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 saving (and, as a result, 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 borrow and hold equity.

As one would expect in a life-cycle model, borrowing (saving) and equity holding are highly sensitive to the shape of the age-income profile. Panels 1 through 5 of Table 8 show consumption, stock and bond holdings, debt, and social security wealth for household heads age 22 to 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. Figure 1 shows the age-income profiles for different educational attainments as estimated by Gourinchas and Parker. Panels 2 to 4 show the estimates of our model with these alternative age-income profiles. High-school-educated households (Panel 2) borrow significantly less than the baseline; college-educated households (Panel 3) borrow significantly more. Notice from Figure 1 that the age-income profile is less steep for high-school-educated households, particularly before age 30, compared with the baseline. In contrast, Figure 1 also shows that the age-income profile for college-educated households is much steeper than the baseline until age 30. One interesting thing to note is that

college-educated households accumulate debt only from ages 22 to 24 – after that they decumulate. With a flat income profile (Panel 4), households never borrow and by the age 37-to-39 period, they accumulate assets worth nearly three times annual income; households in the baseline specification accumulate none by the age of 39. Panel 5 shows that increasing the subjective discount factor reduces debt and increases asset accumulation, as one would expect.

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 matching that of the 25th-percentile household in Table 4. The household in the model starts a little lower, accumulates more, and has its debt level peak 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.

Lastly, 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 that can trade equity compared with one that cannot. For our specifications with realistic age-income profiles, the ability to invest in equity increases lifetime consumption by 2 to 3 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 in our model? The reason is simple. In the Merton-Samuelson model, households can borrow at the return on a riskless bond (that is, 2 percent in this model). Households can realize arbitrarily large returns with no investment by borrowing at 2 percent and investing in equity with an expected return of 8 percent. Crucially, they can also use low-interest unsecured borrowing to borrow against these future 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. This is because household income is low while one is young and it is expensive for households to borrow. Such a model with a wedge between borrowing and lending rates predicts equity holdings that are more in line with household data.

3.4 Effect of social security over the life-cycle for disciplined households

Our simulations show that for disciplined households the existing social security system leads to significantly higher levels of debt, lower private savings, lower consumption, and lower utility. We first discuss life-cycle aspects of social security and then consider overall measures. The deleterious effect of social security for disciplined households follows directly from our assumptions about borrowing. 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 in equities, since an unconstrained household would have invested the optimal amount in equity anyway.

Figure 2 shows results of our simulations with our baseline social security system. What are the differences between the choices disciplined households face with and without social security?

First, the upper left panel of Figure 2 shows that the household with social security accumulates significantly more debt. Initially, the differences are quite small; at age 26, social security leads to a small increase in debt. But at age 30, the household without social security stops accumulating additional debt whereas the household facing social security continues. The contrast is most stark at age 36: The household with social security has more than $8,000 of debt and the same household without has none.

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, and that increases saving. At the same time, households are liberated from the internal rate of return for a large portion of their investments. The upper left panel of Figure 3 illustrates this point. 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 40 and never exceeds 50 percent. The line marked "optimal" shows the percentage for a household that faces no social security system – which equals 100 percent except for a couple of years shortly before retirement. The change in portfolio returns tilts total household income and induces borrowing early in life.

Second, social security depresses stock and bond holding throughout the life cycle but increases total saving early in the life cycle. The reduction in stock and bond holding under a system with social security follows from the increased borrowing documented above. Households are forced to save when they do not want to. As a result, they borrow to increase consumption. In our model, households that borrow

never invest in bonds or stocks because the cost of borrowing always exceeds the returns on stocks and bonds. Thus, the wealth accumulation phase of the life cycle starts only when the debt phase ends – and social security pushes that age back from 36 to 41. (See the upper right panel of Figure 2.) But if we include social security wealth, the picture is quite different. Forced social security saving exceeds debt from the beginning, which means that total wealth for households with social security is always positive. (See the lower left panel of Figure 2.) By contrast, households without social security have negative total wealth until they emerge from debt at age 35.

Third, social security depresses consumption throughout the life cycle. As the lower right panel of Figure 2 shows, the social-security-induced consumption gap is roughly constant at a little more than $1000 until shortly before retirement. Throughout working life, social security drives down after-tax income. Households can (and do) borrow to make up for the forced social security saving. But, initially, households without social security borrow comparable amounts and don't pay social security, yielding higher overall consumption. When households without social security start to accumulate wealth at age 35, the gap actually narrows slightly and then widens shortly before retirement as the effects of higher investment in equities, described above, kick in.

Overall, social security depresses consumption, increases debt, decreases saving, and decreases utility for disciplined households. Table 9 shows that on a population-weighted basis, households with social security (Panel 1) consume $2,000 less per year on average, have $1,600 more in debt, and save about $40,000 less than otherwise-identical households that face no social security (Panel 11). Eliminating social security raises lifetime certain-equivalent consumption by 7.5 percent.

3.5 "Tempted" households

We now show that if tempted households cannot borrow, then they set consumption equal to income in every period and thus behave in the same way as the "rule-of-thumb" consumers of Campbell and Mankiw (1989). The logic follows from the fact that tempted households consume all their available resources each period. In period 0, households have no available resources other than their labor income, so $c = y$, but this implies that $x' = 0$, and therefore $c = y$ in the subsequent period, and so on. Since households always succumb to temptation, the two terms multiplied by ϕ

in equation (1) cancel out, which means that

$$W(0, y) = \sum_{t=0}^{\infty} \delta^t u(y^t).$$

In other words, lifetime utility equals the utility of consuming one's income. Note that while a model with quasi-hyperbolic discounting with $\beta = 1/\infty$ would generate identical behavior, such a model has completely different welfare implications. In the hyperbolic model, a household at time zero would not care at all about the future and would have only a vanishingly small taste for commitment. By contrast, our tempted households care about the future just as much as their disciplined counterparts, differing only in their ability to exercise discipline.

It may seem odd in a paper about borrowing to assume that "tempted" households cannot borrow, but we justify our choice in two ways. First, our data show that borrowing actually occurs mostly in the middle of the income distribution, as lenders mostly exclude low-income households from credit markets. Second, tempted households would simply borrow the maximum possible in their first year of life and then follow our "rule-of-thumb" until they are somehow forced to repay their loans. For some households, such behavior may be realistic, but to analyze properly the effects of our policies, we would need to conjecture about how lenders would adjust credit limits in response to a change in the life-cycle profile of after-tax income.

Tempted households have a pathological taste for commitment. A social planner could use existing assets to impose an income profile equal to the consumption of a disciplined person and thus provide a level of utility equal to that of the disciplined household. Such a policy is unrealistic, but the top panel of Figure 4 shows that for a tempted household, social security provides a reasonable approximation. The dashed line marked "income" shows pre-social-security income. If there were no social security, a tempted household would consume nothing in retirement. With our preference specification, such a household would have infinitely negative utility – certain-equivalent consumption would be zero. Social security reduces income during a household's working life and increases it in retirement, as shown by the line marked "post-SS income." The line marked "optimal" shows the optimal consumption of a baseline, forward-looking household which does not participate in social security. For clarity, we assume that this disciplined household does not trade equity either. Clearly, social security generates a much better consumption stream for the tempted household. But it still has problems. Early in life, consumption is too low; in middle age, it is too high; and in retirement, it is again too low.

4 Alternate social security policy proposals

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. We call this the *payoff* proposal. Second, we consider the whole life cycle and try to design a social security system that doesn't force households to borrow and save at the same time. Specifically, we evaluate a proposal to exempt young households from contributing to social security. We call this the *exemption* proposal.

We evaluate policy proposals using four criteria. First, does the policy proposal prevent disciplined households from simultaneously borrowing and holding social security wealth – the problem which motivated this paper? More generally, what happens to household consumption, saving, and debt? Second, what are the welfare consequences for disciplined and tempted 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 payoff proposal

According to Table 2, our average debtor household is around age 30 and has average debt about 60 percent of income. Suppose we take such a household and assume that it has social security wealth equal to its debt (in reality, its social security wealth is higher). Consider the following policy proposal. Allow that 30-year-old household to withdraw cash from social security. 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.[20] This proposal is not relevant for tempted households since in our model such households do not accumulate debt.

How well does this payoff proposal do? First, does the payoff proposal eliminate the problem of simultaneous unsecured borrowing and social security investment among disciplined households? The answer is: not completely. The upper left panel of Figure 5 shows debt and social security wealth with our proposal ("reduced debt")

[20] 12.8 percent would be the tax rate necessary to pay for the withdrawal of 60 percent of income from accumulated social security wealth to pay off existing debt.

and without ("existing"). Households use most but not all of their social security wealth to pay down debt. A small portion is used to finance increased consumption. 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 save in equities earlier (and reduce debt payments) resulting in a wealth effect, which in turn causes households to want higher consumption.

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

Third, the payoff proposal compares well with alternative proposals for social security. Panels 3 and 4 of Table 10 show the effects on utility of increasing the internal lending rate but otherwise keeping the social security system as is. Allowing households to use social security to pay off debt yields a welfare payoff equivalent to an increase in the internal lending rate (ILR) of 3 to 4 percent for a disciplined household. For a tempted household, an increase in the ILR generates a much bigger increase in utility than our proposal. The difference in effects for different types of households results from the unwillingness of tempted households to trade equity on their own. Disciplined households already buy some equity on their own, so their effective "lending rate" is much higher than the return on social security. But for tempted households, increasing the ILR is akin to forcing them to invest in equity with consequently large welfare improvements. We will return to this theme in the next section when we consider our proposal for exemptions to social security contributions.

To help illustrate the payoff proposal, we 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? One way to view it is that the household is borrowing from

itself. Later on, we will argue that the sensible interpretation is that it is borrowing 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. The 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 The exemption proposal

In this section, we propose changes in the age structure of social security contributions. 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. 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. We consider three age-based exemptions: no contributions before age 30; no contributions before age 40; and no contributions before age 50. As we discuss in the conclusion, 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, but we abstract from these 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.

Again we return to the four questions posed in the beginning of this section. First, do these policy proposals prevent households from simultaneously borrowing and holding social security wealth? Yes; Figure 6 shows that 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. The upper left panel of Figure 2 shows that with the current social security system, households pay off most of their debt by age 40 anyway. But even with the age 30 exemption, households are out of debt by the time social security taxes take effect. Disciplined households anticipate the 12.9 percent permanent drop in income that will occur at age 30 and start saving in advance.

More broadly, the payoff proposal improves the general household balance sheet. Table 9 shows per capita levels of consumption, debt, and savings for disciplined households under the different exemption plans (Panels 2, 3, and 4, respectively). Panel 1 of Table 9 shows the baseline model with the existing social security system. Panel 11 shows the results of the model with no social security system. The age-30

exemption leads to an almost complete elimination of unsecured debt. Households increase consumption and save slightly more. Why do both consumption and saving go up? As we explained in Section 3.1, an exemption as we define it can be interpreted as a loan: 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. Delaying contributions to social security until later in the life cycle leads to continued increases in saving and consumption but also to increases in debt. Why do both debt and saving go up? The longer the exemption, the steeper the income profile when young. (See the bottom panel of Figure 4.) 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 causes a household to increase saving today.

Second, what happens to utility? The exemption generates substantial increases in utility for both disciplined and tempted households. The age-30 exemption increases certainty-equivalent consumption by roughly comparable amounts for disciplined and tempted households – 3.4 and 3.3 percent, respectively. An age-40 exemption helps both types of households – but the gains are bigger for the disciplined households – 5.5 percent versus 4.8 percent. And an age-50 exemption generates a large welfare gain for disciplined households – 6.7 percent of annual consumption. In contrast, the age-50 exemption results in a small reduction in welfare for the tempted household vis-a-vis the age-40 exemption. For the disciplined household, the age-50 exemption reduces the burden of social security by about 90 percent (6.7 percent change in utility/7.5 percent change in utility). In the next section, we explain why the exemptions work so well.

Third, how do the utility gains from the payoff proposal compare with gains from simply raising the internal lending rate? For both disciplined and tempted 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. Then, for disciplined households, exemptions clearly dominate. (See Panels 5, 6, and 7 of Table 9.) Even an increase in the internal lending rate to 5 percent increases certain-equivalent consumption by an amount less than that for the age-30 exemption. By contrast, for tempted households, raising the internal lending rate is a winning proposition. Raising the ILR to 3 percent generates gains comparable to those of the exemptions, but higher ILRs lead to much larger welfare gains. (Compare Panel 4 of Table 9 with Panel 7 of Table 9. This welfare gain for tempted 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 leads to a replacement rate of 88 percent of pre-retirement income, holding contributions constant. In that case, the decline in cash flow at the time of retirement for tempted households is small. Stabilizing income during retirement has large welfare benefits for tempted households.

Suppose we simultaneously lower withholding rates when we increase the ILR so that the replacement rate in retirement remains constant at 43 percent. Now, the results are reversed. The gains to tempted households are relatively small. An increase of the ILR to 5 percent generates a welfare increase in between those generated by the age-30 and the age-40 exemptions. For disciplined households, however, the gains are much larger: An increase in the ILR to 5 percent yields welfare increases comparable to the age-40 exemption. Why are the effects so different? For tempted households, the opportunity to invest at 5 percent is extremely valuable – since they don't invest in equity (or in anything else). But for disciplined households, even a 5 percent ILR is still less than they get investing in equity 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. So we ask how the payoff proposal (and raising the ILR) would affect the level of net contributions. To answer this question, we assume that aggregate wage growth equals the baseline ILR of 2 percent. This means that if there were no population growth, social security would always take in exactly as much in contributions as it pays out in benefits. In addition, our exemptions should have no effect on net contributions. In fact, given the year-2000 population distribution, our exemptions do affect net contributions. The age-40-to-age-50 group is relatively large and so the age-30 and age-40 exemptions shift contributions to that group, which raises net contributions. The age-50 exemption shifts contributions away from the age-40-to-50 group and reduces net contributions. In contrast, raising the ILR always leads to a deterioration in the level of net contributions. However, proposals to raise the ILR are always paired with proposals to generate additional income through investment in stock. But our results on net contributions illustrate that we can generate significant welfare improvements without investing in stock or reducing the solvency of the social security system. In Section 4.4, we show that by changing the internal borrowing rate we improve the solvency of social security while preserving

some of the welfare benefits.

4.3 Approximating optimal policy rules

Let's first consider the tempted households. The top panel of Figure 4 shows that the existing social security system lowers consumption to below the optimal level early in life, allows it to exceed the optimal level in middle age, and allows it to fall short of the optimal level in retirement. The bottom panel of Figure 4 shows after-social-security income for the three exemptions compared with optimal consumption for a household that can't trade equity. All three exemptions increase consumption early in life and lower it in middle age. In each case, the cost of the exemption is a discontinuous fall in consumption at the start of the exemption. The cost of the discontinuity illustrates why tempted households prefer the age-40 to the age-50 exemption in spite of the fact that from age 50 to retirement the age-50 exemption almost exactly matches the optimal profile.

For the disciplined households, the explanation for the success of the exemptions is more complex because social security interferes not only with saving decisions but also with portfolio allocation decisions. We first consider the saving decision and then the portfolio allocation decision. The top right panel of Figure 3 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 figure shows, deviations from the optimum get progressively smaller as we extend the exemption.

We note two odd features of the excess saving. First, households in all the social security systems oversave relative to the no-social-security benchmark until roughly age 55. For households in the existing social security system, this is not very 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 these households oversave even before the exemption starts? The reason is that the exemption leads to a discontinuous but predictable drop in income (17.9 percent for the age-40 exemption and 30.4 percent for the age-50 exemption; see Table 9). Households engage in private saving in anticipation of this drop. See the lower panel of Figure 3, 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? The lower panel of Figure 3 shows that private saving falls relative to the no-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. The upper left panel of 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 the figure. Households invest almost all their money in equity in virtually all portfolio choice models with labor income; this 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 50 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 percent in our case.

The top left panel of Figure 3 shows that exemptions close the gap. For the age-40 and age-50 exemptions, households never invest less than 50 percent in equities, and they invest 100 percent 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. It starts saving before the head reaches age 30 in anticipation of a decline in income when the exemption ends. At age 30, it starts dissaving to maintain consumption, but by the age of 35, it is accumulating again, and all its accumulation goes into equity. Now consider a household facing an age-40 or age-50 exemption. In both cases, the household saves entirely in equities in anticipation of the exemption. Then, when the exemption ends, its starts to draw down its private saving and replace it with forced saving, which it cannot invest in equities.

To quantify the effect of exemptions on optimal saving and on optimal portfolio allocation, we conduct the following exercise. Table 11 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 $1,000 compared with the $1,700 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 eliminates only 45 percent of the portfolio loss

and an age-40 exemption eliminates about 70 percent. As the top left panel of Figure 3 shows, going from the age-40 exemption to the age-50 exemption brings portfolio allocation much closer to the optimal level.

4.4 Raising the internal borrowing rate

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 two sections, implicitly 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 2 percent?

In Table 12, 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.

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 consider two policy experiments aimed at resolving this problem in the context of a life-cycle model with both disciplined and tempted households. First, we consider allowing households to use the money in social security to get out of debt. And then we consider options to change the age structure of social security to prevent households from borrowing while they also contribute to social security. We find that both proposals, but in particular the latter, solve the problem in question and lead to significant increases in household welfare, consumption, and saving and to reductions in high-interest unsecured debt.

We show that our options generate 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: It does not subject tempted households to politically unacceptable risks.

A model that incorporates both endogenous labor supply and general equilibrium would strengthen our results significantly. However, 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 – 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 9 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(5): 1079-97.

[2] Akerlof, George. 1998. Self-Control and Saving for Retirement: Comment. *Brookings Papers on Economic Activity.* 0(1)185-189.

[3] Altig, David and Steven J. Davis. 1992. The timing of intergenerational transfers, tax policy and aggregate saving. *American Economic Review,* 82(5):1199-1220.

[4] 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(4):832-857.

[5] Bisin, Alberto and Piero Gottardi. 1999. Competitive equilibria with asymmetric information. *Journal of Economic Theory* 87(1):1-48.

[6] Bohn, Henning. 1999. Should the social security trust fund invest in equities? An intergenerational welfare analysis. *Review of Economic Dynamics* 2(3):666-697.

[7] 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.

[8] Campbell, John Y. and Luis Viceira. 2002. *Strategic Asset Allocation.* New York: Oxford University Press.

[9] Campbell, John Y., Joao Cocco, Francico 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.

[10] 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.

[11] 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.

[12] Davis, Steven J., Felix Kubler, and Paul S. Willen. 2003. Borrowing costs and the demand for equity over the life-cycle. Revised version of NBER working paper 9331 available at `http://gsbwww.uchicago.edu/fac/paul.willen/research/`.

[13] Diamond, Peter A. 1977. A framework for social security analysis. *Journal of Public Economics* 8(3):275-98.

[14] Diamond, Peter A. and John Geanakoplos. 2003. Social security investment in equities. *American Economic Review* 93(4):1047-1074.

[15] Dubey, Pradeep, John Geanakoplos, and Martin Shubik. 2003. Default and punishment in general equilibrium. Yale University. Working paper.

[16] Durkin, Thomas A. 2000. Credit cards: Use and consumer attitudes. *Federal Reserve Bulletin* (September)86(9):623-634.

[17] 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(3):303-26.

[18] Feldstein, Martin. 1985. The optimal level of social security benefits. *Quarterly Journal of Economics* 100(2):303-320.

[19] 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.

[20] Feldstein, Martin and Elena Ranguelova. 2001. Individual risk in an investment-based social security system. *American Economic Review* 91(4):1116-1125.

[21] 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.

[22] Gokhale, Jagadeesh, Laurence J. Kotlikoff, and Todd Neumann. 2001. Does participating in a 401(k) raise your lifetime taxes? NBER Working paper 8341.

[23] Gomes, Francisco and Alex Michaelides. 2002. Life cycle asset allocation: A model with borrowing constraints, uninsurable labor income risk and stock-market participation costs. London Business School. Working Paper.

[24] Gourinchas Pierre-Olivier and Jonathan A. Parker. 2002. Consumption over the life cycle. *Econometrica* 70(1):47-89.

[25] 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(1):149-85.

[26] Gul, Faruk and Wolfgang Pesendorfer. 2001. Temptation and self-control. *Econometrica* 69(6):1403-35.

[27] Gul, Faruk and Wolfgang Pesendorfer. 2004a. Self control and the theory of consumption. *Econometrica* 72(1) 119-158.

[28] Gul, Faruk and Wolfgang Pesendorfer. 2004b. Self control, revealed preference and consumption choice. *Review of Economic Dynamics* 7(2):243-264.

[29] 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.

[30] Hurst, Erik, Ming Ching Luoh, and Frank Stafford. 1998. The Wealth Dynamics of American Families, 1984-1994. *Brookings Papers on Economic Activity* 0(1) 267-329.

[31] Hurst, Erik. 2003. Grasshoppers, ants and pre-retirement wealth: A test of permanent income consumers. University of Chicago working paper.

[32] Imrohoroglu, Ayse, Selahattin Imrohoroglu, and Douglas Joines. 1995. A life-cycle analysis of social security. *Economic Theory* 6(1):83-114.

[33] Imrohoroglu, Ayse, Selahattin Imrohoroglu, and Douglas Joines. 2003. Time inconsistent preferences and social security. *Quarterly Journal of Economics* 118(2):745-784.

[34] Judd, Kenneth, Felix Kubler, and Karl Schmedders. 2003. 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.

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

[36] 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.

[37] Krusell, Per and Anthony Smith. 2003. Consumption-Savings Decisions with Quasi-geometric Discounting. *Econometrica* 71(1) 365-375.

[38] Laibson, David. 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112(2) 443-477.

[39] Laibson, David, Andrea Repetto, and Jeremy Tobacman. 1998. Self-control and saving for retirement. *Brookings Papers on Economic Activity* 1:91-172.

[40] 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.

[41] Scholz, John Karl, Ananth Seshadri, and Surachai Khtiatrakun. 2003. Are Americans saving optimally for retirement? University of Wisconsin - Madison working paper.

[42] Smetters, Kent and Jan Walliser. 2004. Opting out of Social Security. *Journal of Public Economics*. 88(7-8) 1295-1306.

[43] 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: Summary of results. Certain-equivalent lifetime consumption measured in thousands of 1987 dollars. Numbers in parentheses are percentage difference from existing.

	"Disciplined" households	**"Tempted" households**
Existing	22.2 (-0.0%) Suboptimal portfolio choice	20.5 (0.0%) Reasonable pension
No social security	23.9 (7.5%) Optimal portfolio choice	0.0 (-100.0%) Starvation
Age 30 exemption	23.0 (3.4%)	21.1 (3.3%)
Age 40 exemption	23.4 (5.5%)	21.4 (4.8%)
Age 50 exemption	23.7 (6.7%)	21.4 (4.6%)
Raise IRR to 4%	22.5 (1.4%)	21.9 (7.0%)
Age 40 exemption + Raise IBR to 5%	23.0 (3.4%)	21.1 (3.0%)

Note:

1. See Section 3 for a complete description of the model.

2. See equation (1) for the calculation of certain-equivalent (CE) consumption.

3. With no social security, tempted households earn no income in retirement and thus consume nothing. Since utility is Constant Relative Risk Aversion with Relative Risk Aversion equal to 3, certain-equivalent consumption is zero.

Table 2: Descriptive statistics for 1999 PSID sample, by debt holding.

		I	*II*	*III*	*IV*
		Sample by household unsecured debt level			
Variable		All	No debt	Positive debt	p-value of difference
Age of the household head		31.6	32.3	31.1	<0.01
Percent of households with...					
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
Unsecured 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
Quantities divided by income:					
Head's total SS contributions	Mean	0.85	0.91	0.81	<0.01
	Median	0.64	0.68	0.62	0.02
Household unsecured 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	

Note: Sample includes 1) all unmarried households with heads between the ages of 22 and 39, and 2) married households with heads between the ages of 22 and 39 where the head also has a complete earnings history. Head's total social security contributions are computed 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. Household net liquid assets are defined as the sum of cash, checking and saving balances, stocks, corporate bonds, and Treasuries less unsecured debt. Household net financial assets are defined as the sum of stocks and corporate bonds less unsecured debt. Sample is split by households with no unsecured debt (column II) and households with some positive amount of unsecured debt (column III). All dollar amounts are reported in 1996 dollars.

Table 3: Portfolio composition of PSID households under 40, by debt holding. "Cash" includes cash and checking and saving accounts. See note to Table 2.

Variable	*I*	*II*	*III*	*IV*
		Sample by household debt level		
	All	No debt	Positive debt	p-value of difference between II and III
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 stocks or bonds	0.32	0.30	0.33	0.26
Owning cash, stocks or bonds	0.85	0.77	0.90	<0.01
Owning a business	0.09	0.10	0.09	0.91
Having negative net worth	0.16	0.01	0.26	<0.01
Having negative net liquid assets	0.39	0	0.63	<0.01
Having negative net financial assets	0.50	0	0.82	<0.01
Sample Size	2,077	844	1,233	

Table 4: Median ratio of net financial assets to household income, by age range. See note to Table 2.

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 5: Distribution of ratio of unsecured debt to income before and after adjusting for social security wealth.

	I All households	*II* All households with debt	*III* Single households with debt
Panel A			
Percent with positive unsecured debt			
	0.62	1.00	1.00
Distribution of debt-to-income ratio			
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
Panel B			
Percent with positive social security augmented debt			
	0.17	0.28	0.31
Distribution of social security augmented debt-to-income ratio			
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

Note: Data are from the 1999 PSID. Sample in columns I and II includes 1) all unmarried households with heads between the ages of 22 and 39, and 2) married households with heads between the ages of 22 and 39 where the head also has a complete earnings history (2,077 households). Sample in column IV includes all unmarried households with heads between the ages of 22 and 39 (513 households). Social security augmented debt is defined as unsecured 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 unsecured debt, social security augmented debt is set equal to zero.

Table 6: Aggregate effect of allowing households to use social security wealth to pay off non-collateralized debt (for households age 22 to 40 in 1999).

Average debt reduction for households with debt	$6,225
Average annual interest saving resulting from debt reduction for households with debt	$498
Total debt reduction for households with bonds	$144.1 billion
Total annual interest saved by households with bonds	$11.5 billion

Note: All dollar amounts are in 1996 dollars. Data are 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%, while social security wealth assumes an interest rate of 2%. To compute debt reduction for the total U.S. population, we use 104 million households (Census Bureau), 36.6% of which have heads between the ages of 22 and 40 (PSID weighted data) and 60.8% of which have non-collateralized debt (PSID weighted data).

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%	
$\mathrm{E}(\tilde{r}_E)$	8%	
$\mathrm{std}\,(\tilde{r}_E)$	15%	
Social security parameters		
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: Consumption and portfolio choice for disciplined households over the life cycle as predicted by the model under various assumptions.

#	β	Profile	Equity benefit		Age of Household Head						
					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

Note:

1. All parameters are baseline (Table 7) unless otherwise indicated.
2. Consumption is measured in 1987 dollars.
3. Equity benefit measures the percentage increase in certain-equivalent consumption yielded by investment in equity.

Table 9: Effects of exemptions. All amounts are in thousands of 1987 dollars unless otherwise noted. See Table 10 for definitions.

#	Parameters of SS system: Start age	Parameters of SS system: With-holding	Parameters of SS system: ILR	Parameters of SS system: Repl. rate	Utility: Disciplined (Δ)	Utility: Tempted (Δ)	Population weighted *per capita*: Disciplined Cons.	Population weighted *per capita*: Disciplined Debt	Population weighted *per capita*: Disciplined Saving	Population weighted *per capita*: Tempted Cons.	Population weighted *per capita*: Net SS contrib.
1	21	10.6	2	43	22.2 (-0.0)	20.5 (0.0)	26.1	3.4	58.3	25.5	1.1
2	30	12.9	2	43	23.0 (3.4)	21.1 (3.3)	26.4	0.2	60.8	25.2	1.3
3	40	17.9	2	43	23.4 (5.5)	21.4 (4.8)	26.8	0.8	72.4	25.2	1.4
4	50	30.4	2	43	23.7 (6.7)	21.4 (4.6)	27.3	1.3	84.9	25.6	1.0
5	21	10.6	3	54	22.3 (0.6)	21.3 (4.0)	26.4	4.0	52.3	25.8	0.7
6	21	10.6	4	68	22.5 (1.4)	21.9 (7.0)	26.7	4.7	45.9	26.2	0.2
7	21	10.6	5	88	22.7 (2.3)	22.3 (9.1)	27.2	6.0	38.9	26.8	-0.5
8	21	8.6	3	43	22.7 (2.2)	20.8 (1.8)	26.7	3.4	60.2	26.0	0.6
9	21	6.9	4	43	23.1 (4.0)	21.1 (3.2)	27.2	3.4	61.9	26.5	0.1
10	21	5.4	5	42	23.4 (5.6)	21.3 (3.9)	27.5	3.5	63.5	26.9	-0.3
11	21	0.0	2	0	23.9 (7.5)	0.0 (-100.0)	28.1	1.8	99.0	26.9	0.0

Table 10: The effects of allowing households to use social security wealth to pay off debt.

#	Wealth at Age 30		Parameters of SS system			Utility	
	Non-SS	SS	With-holding	ILR	Repl. rate	Disciplined (Δ)	Tempted (Δ)
1	-0.6	0.6	10.6	2	42	25.3 (0.0)	21.9 (0.0)
2	0.0	0.0	12.8	2	43	25.9 (2.1)	22.6 (2.9)
3	-0.6	0.6	10.6	3	51	25.6 (1.1)	23.4 (6.9)
4	-0.6	0.6	10.6	4	64	25.9 (2.4)	24.7 (12.5)
5	0.0	0.0	0.0	2	0	27.7 (9.5)	0.0 (-100.0)

Notes:

1. ILR is the internal lending rate (See Section 3.1 for an explanation.)
2. Private saving equals stocks plus bonds minus debt.
3. Scenario is baseline (see Table 7).
4. Replacement rate is the annuity value of the lump-sum payment from social security as a fraction of the top 35 years of income.
5. Utility is certain-equivalent consumption (see equation (1)).
6. Δ is difference from existing.

Table 11: Effects of exemptions. Scenario is baseline (see Table 7) except that households cannot trade equity. See Table 10 for definitions.

	Parameters of SS system				Utility		Population weighted *per capita*				
							Disciplined			Tempted	Net SS
#	Start age	With-holding	ILR	Repl. rate	Disciplined (Δ)	Tempted (Δ)	Cons.	Debt	Saving	Cons.	contrib.
1	21	10.6	2	43	21.7 (0.0)	20.5 (0.0)	24.1	5.6	66.3	25.5	1.1
2	30	12.9	2	43	22.3 (3.1)	21.1 (3.3)	24.4	0.9	65.5	25.2	1.3
3	40	17.9	2	43	22.6 (4.4)	21.4 (4.8)	24.2	2.8	61.9	25.2	1.4
4	21	0.0	2	0	22.7 (4.5)	0.0 (-100.0)	24.2	3.2	77.4	26.9	0.0

Table 12: Effects of changes in the Internal Borrowing Rate (IBR) for an age-40 exemption. (See Section 3.1 for an explanation). Scenario is baseline (see Table 7) except where noted. See Table 10 for definitions.

#	Parameters of SS system				Utility		Population weighted *per capita*					
	Start age	With-holding	IBR	Repl. rate	Disciplined (Δ)	Tempted (Δ)	Disciplined Cons.	Disciplined Debt	Disciplined Saving	Tempted Cons.	Net SS contrib. 2000	Net SS contrib. 2020
1	21	10.6	2	43	22.2 (-0.0)	20.5 (0.0)	26.1	3.4	58.3	25.5	1.1	-0.1
2	40	17.9	2	43	23.4 (5.5)	21.4 (4.8)	26.8	0.8	72.4	25.2	1.4	0.0
3	40	21.6	4	43	23.1 (4.2)	21.2 (3.7)	26.4	0.6	74.5	24.6	2.0	0.6
4	40	24.0	5	43	23.0 (3.4)	21.1 (3.0)	26.1	0.4	76.0	24.2	2.4	1.0
5	40	26.9	6	43	22.7 (2.4)	20.9 (2.0)	25.8	0.3	78.2	23.7	2.9	1.4
6	40	30.3	7	43	22.5 (1.2)	20.6 (0.7)	25.5	0.2	81.0	23.1	3.5	2.0
7	40	34.4	8	43	21.3 (-3.9)	20.3 (-1.0)	25.0	0.2	84.0	22.4	4.2	2.6
8	21	0.0	2	0	23.9 (7.5)	0.0 (-100.0)	28.1	1.8	99.0	26.9	0.0	0.0

Figure 1: Expected labor income profiles.

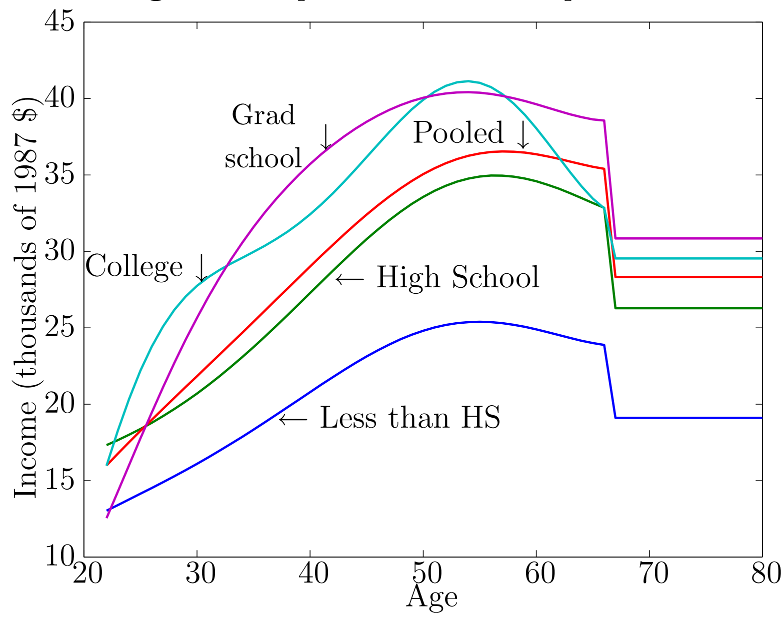

Figure 2: Effect of social security on various household indicators over the life cycle. Baseline specification.

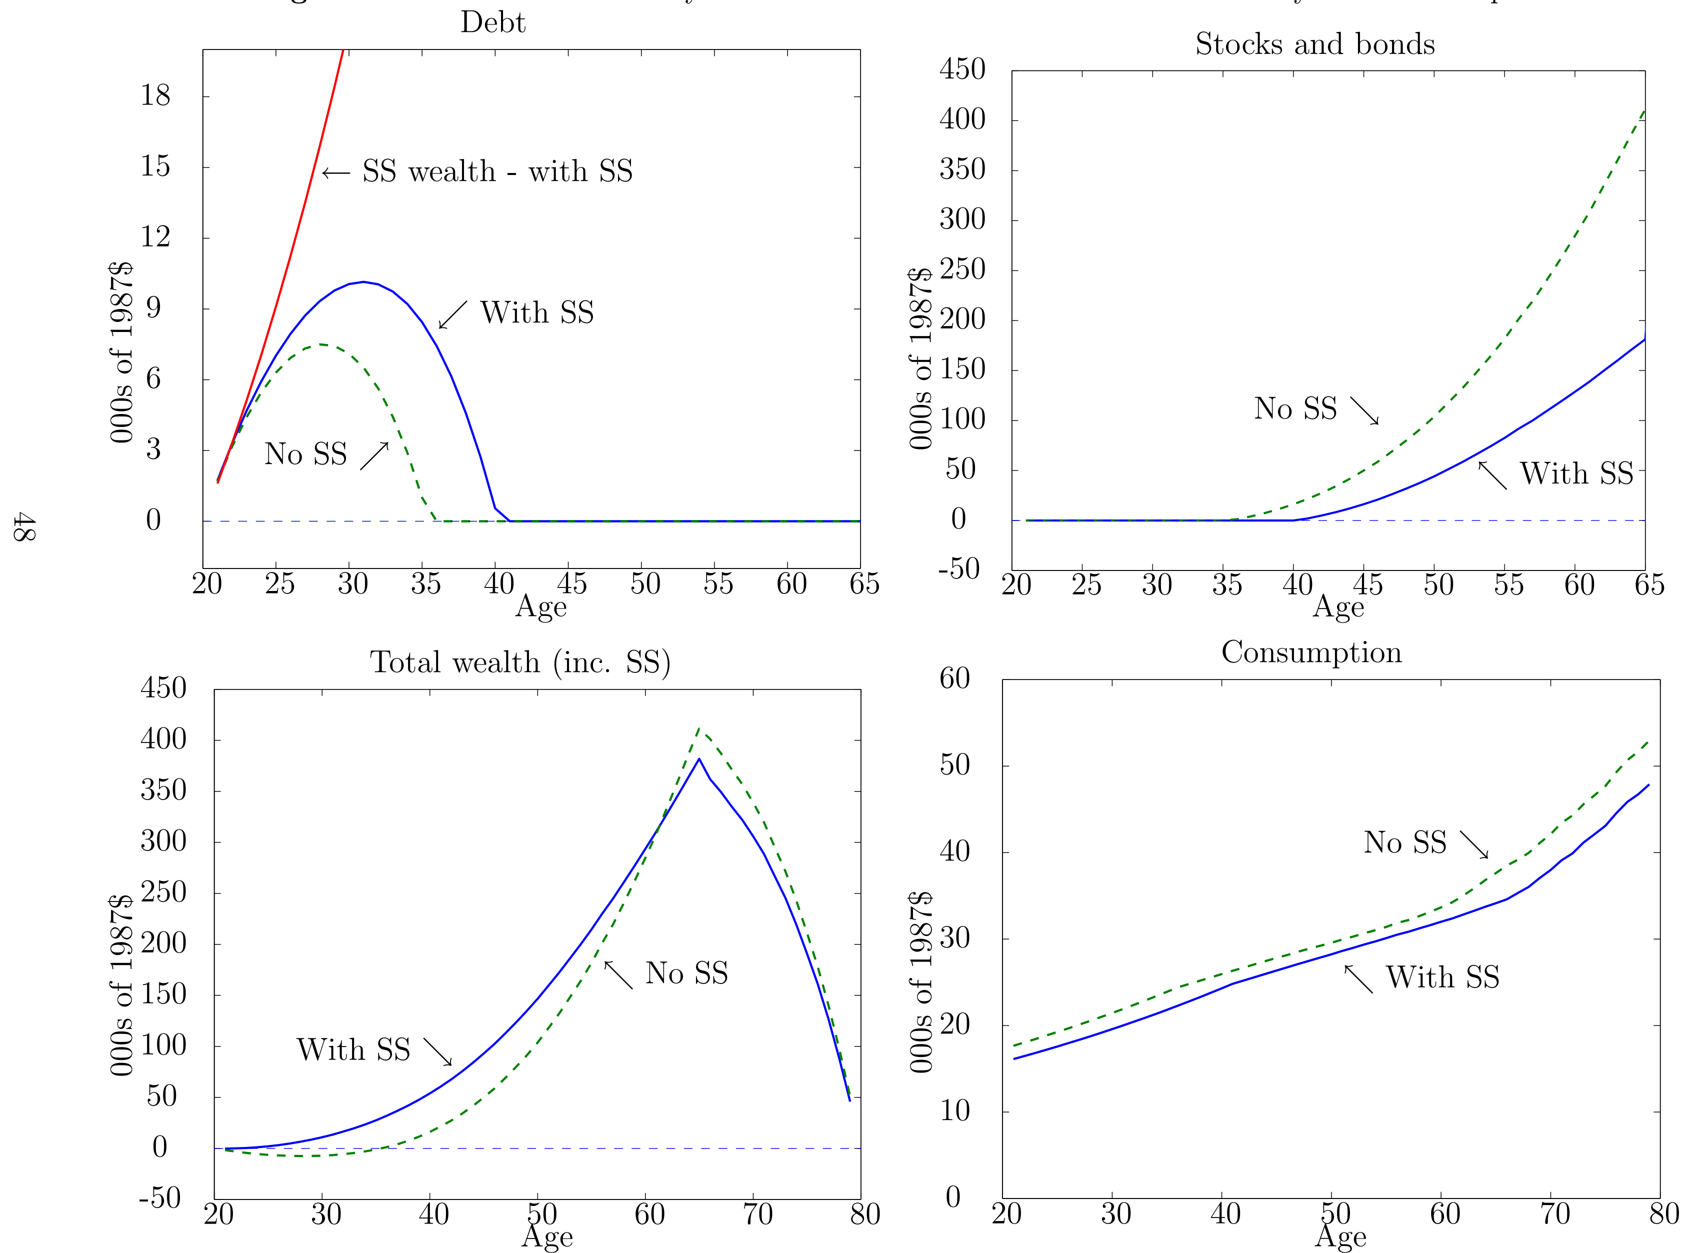

Figure 3: Excess saving over the life cycle for various different social security schemes.

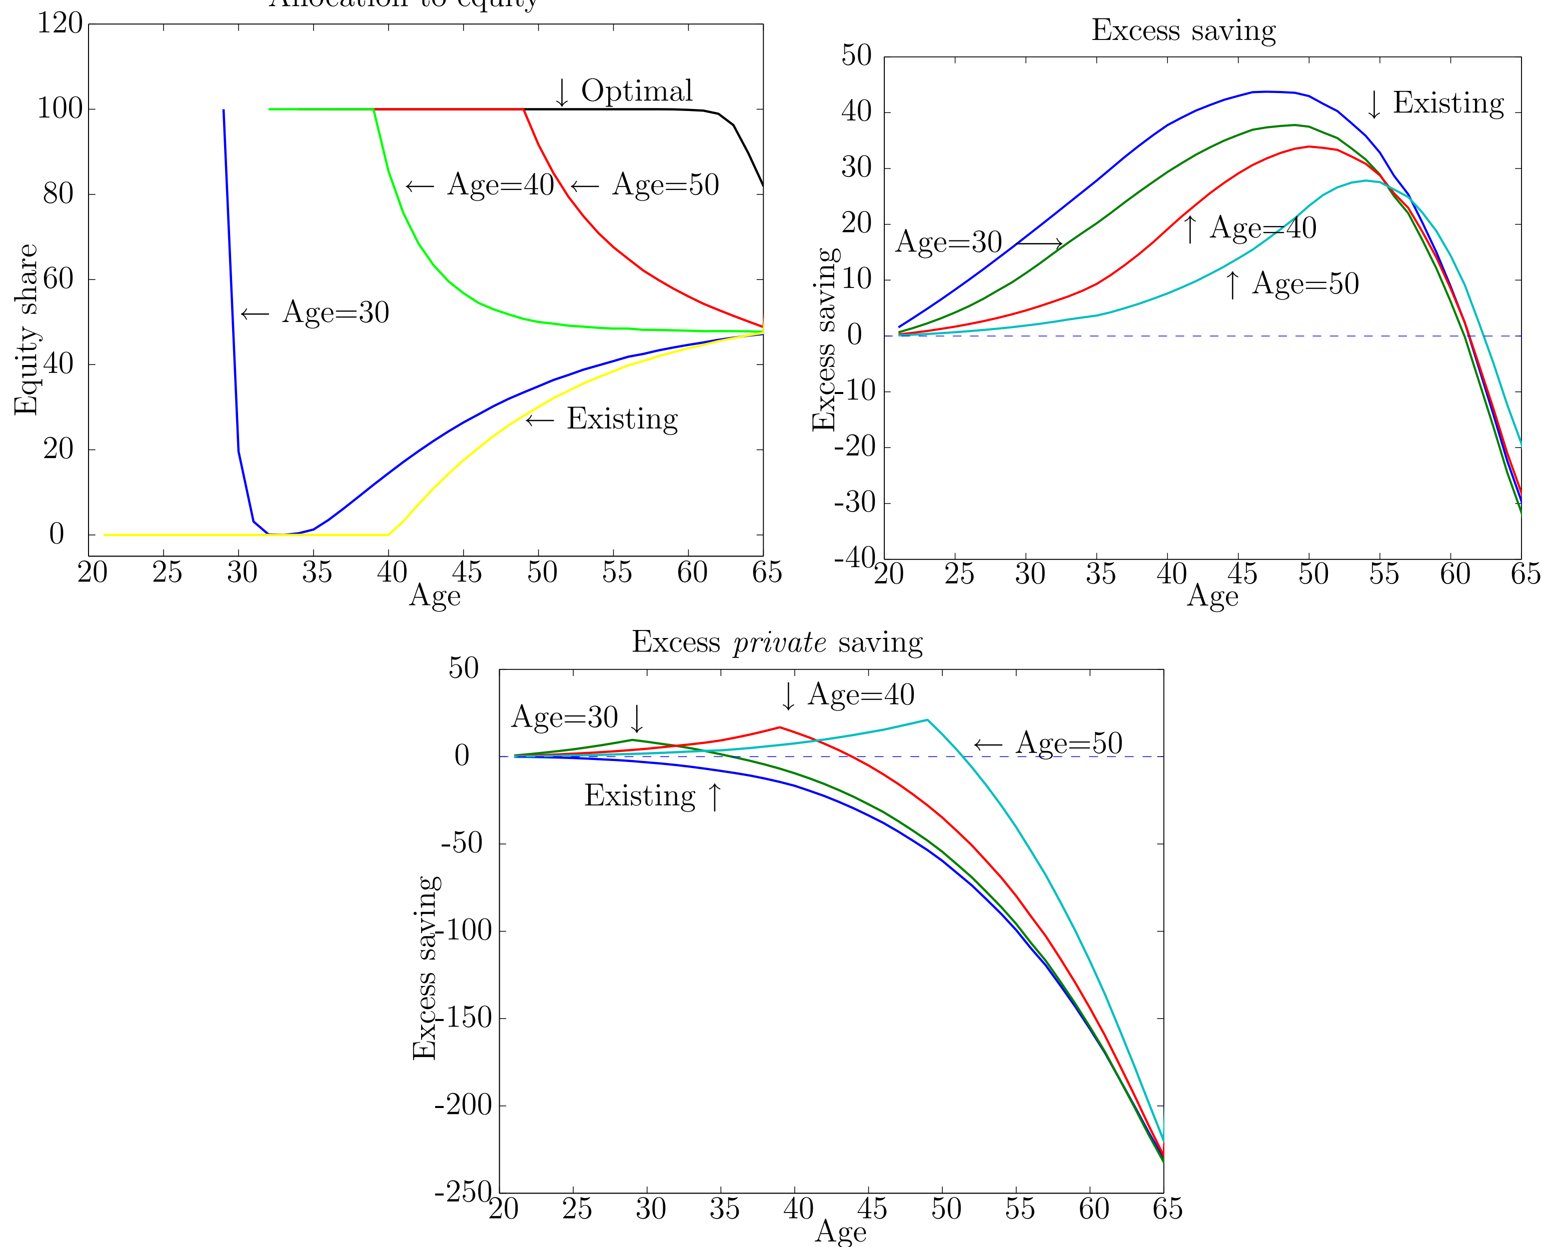

Note: Excess saving equals the difference between total saving (including social security) under the listed scheme and total saving under the optimal scenario for the household in which there is no social security. Age 30 refers to an exemption ending at age 30; and so on. Existing is the existing system.

Figure 4: The effects of exemptions for tempted households.

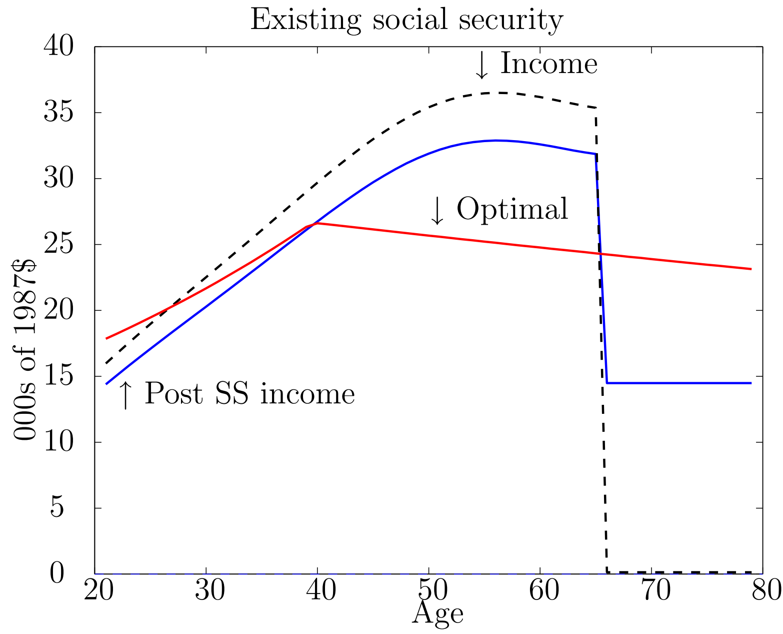

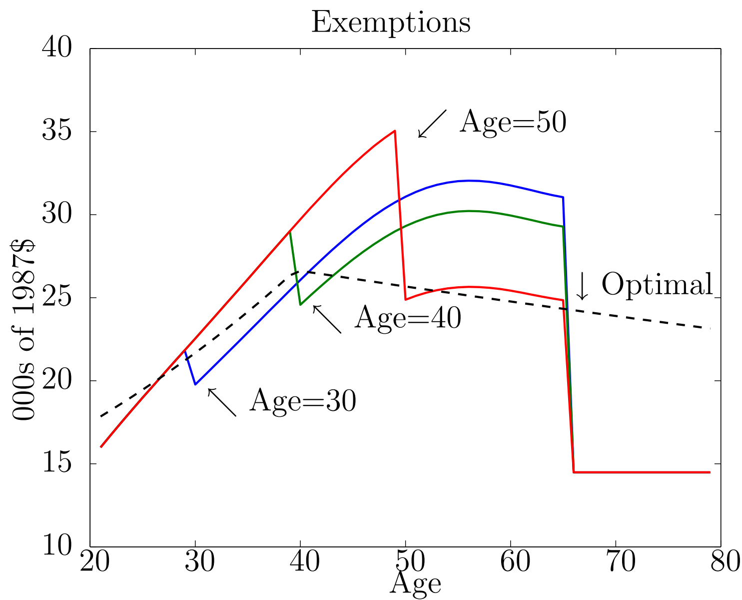

Figure 5: Effects of allowing households to use social security wealth to pay off debt for a household with debt equal to 60 percent of annual income.

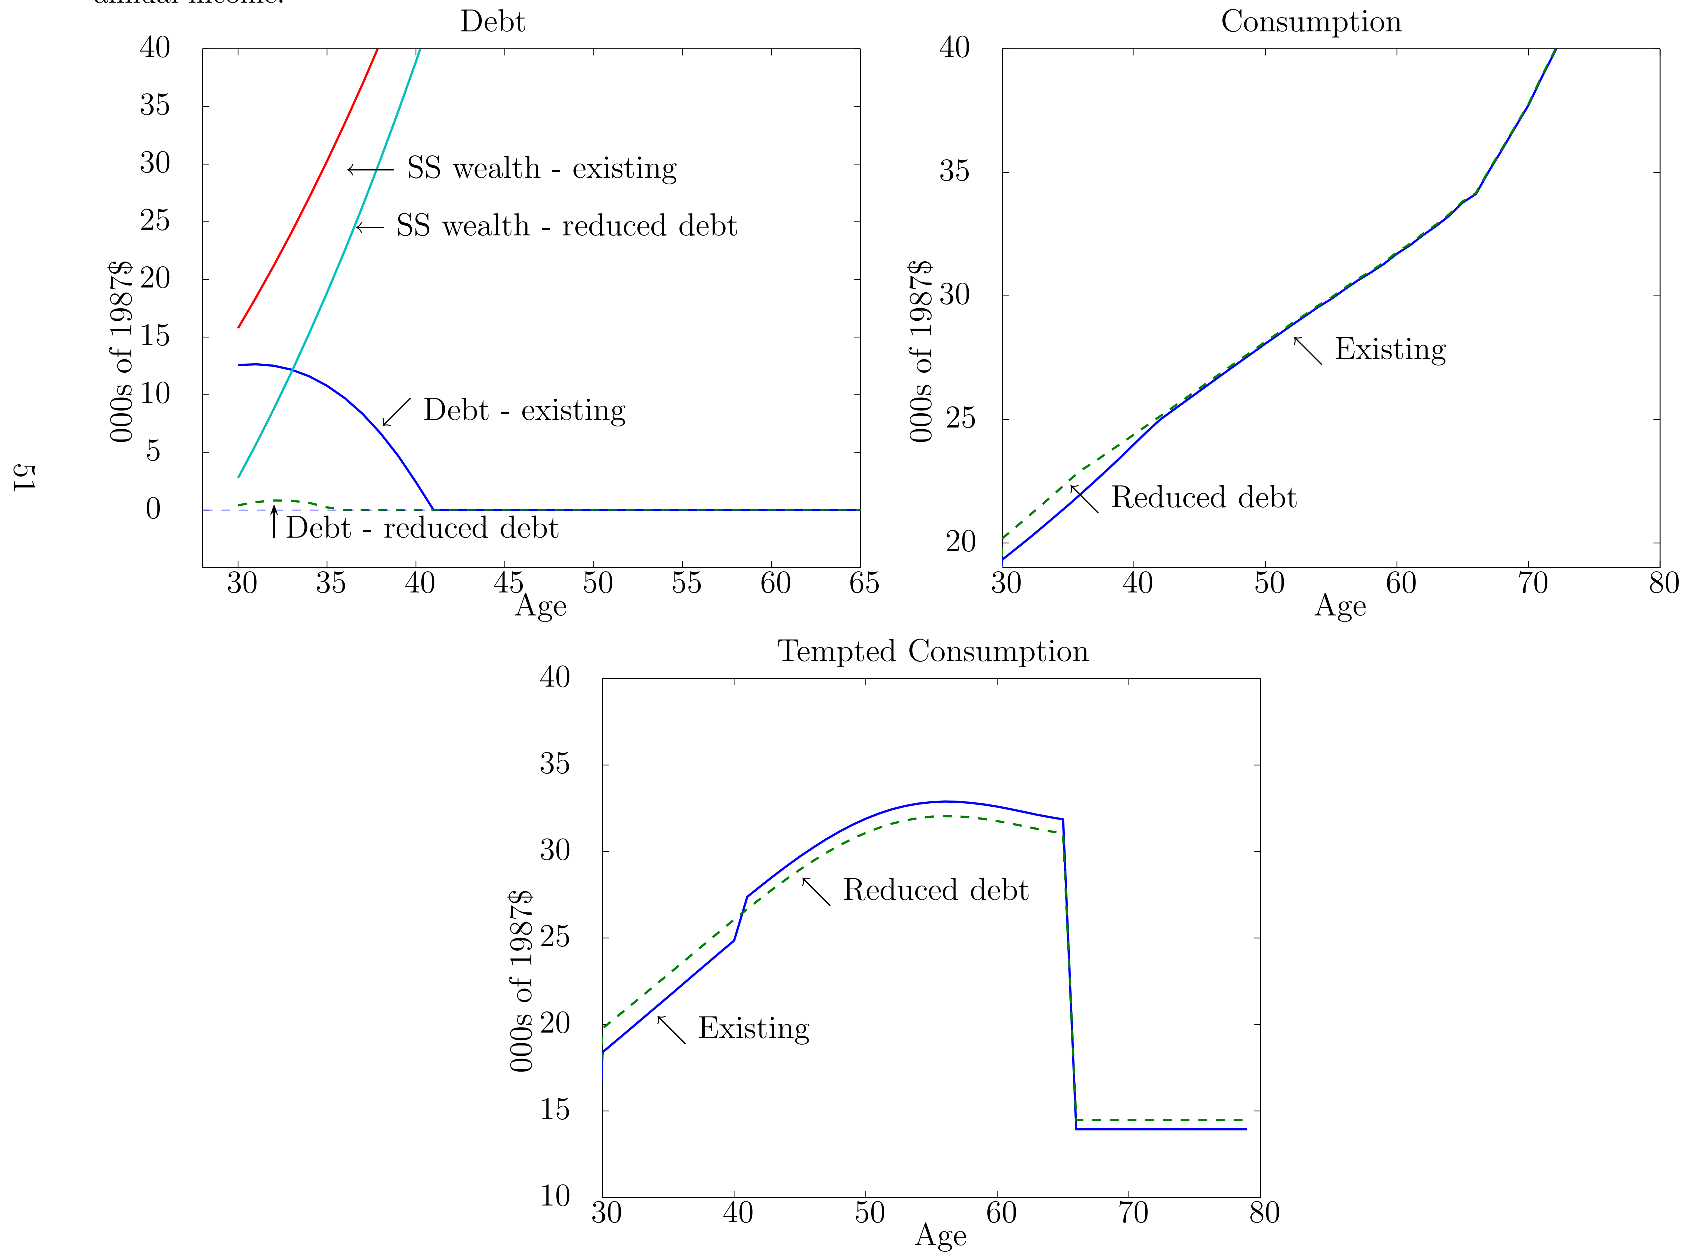

Figure 6: The effect of exemptions on debt and social security wealth.