Borrowing costs and the demand for equity over the life cycle*

Steven J. Davis†
Graduate School of Business
University of Chicago
and NBER

Felix Kubler‡
Department of Economics
Stanford University

Paul Willen§
Graduate School of Business
University of Chicago

October 1, 2002

*We thank seminar participants at the Minnesota Macroeconomics Summer Institute, the NBER Summer Institute, Notre Dame University and the University of Chicago for many helpful comments, Amir Yaron for a thoughtful discussion at the NBER Summer Institute, Jonathan Parker for providing us with the parameters of the income processes in Gourinchas and Parker (2002), Nick Souleles for directing us to data on charge-off rates for consumer loans and David Arnold, Jeremy Nalewaik and Stephanie Curcuru for able research assistance. We gratefully acknowledge research support from the Graduate School of Business at the University of Chicago.

†Phone: (773) 702-7312. E-mail: steve.davis@gsb.uchicago.edu

‡Phone: (650) 724-4902. E-mail: fkubler@stanford.edu.

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

Borrowing costs and the demand for equity over the life cycle

by Steven J. Davis, Felix Kubler and Paul Willen

October 1, 2002

Abstract

We analyze consumption and portfolio behavior in a life-cycle model calibrated to U.S. data on income processes, borrowing costs and returns on risk-free and equity securities. Even a modest wedge between borrowing costs and the risk-free return dramatically shrinks the demand for equity. When the cost of borrowing equals or exceeds the expected return on equity – the relevant case according to the data – households hold little or no equity during much of the life cycle. The model also implies that the correlation between consumption growth and equity returns is low at all ages, and that risk aversion estimates based on the standard excess return formulation of the consumption Euler Equation are greatly upward biased. The bias diminishes, but remains large, for "samples" of households with positive equity holdings. The demand for equity in the model is non-monotonic in borrowing costs and risk aversion. The standard deviation of marginal utility growth is an order of magnitude smaller than the Sharpe ratio.

1 Introduction

We analyze consumption and portfolio behavior in a life cycle model with borrowing costs that exceed the risk-free investment return. We calibrate the model to U.S. data on the cost of borrowing, returns on risk-free and equity securities and life-cycle income processes for several education groups. Consumers in the model exhibit standard time-separable preferences with modest risk aversion and mild impatience.

The wedge between borrowing costs and the risk-free return has several important effects on consumption and portfolio behavior. First, even a modest wedge dramatically shrinks the demand for equity throughout the life cycle. Second, when the borrowing rate equals or exceeds the expected return on equity – the relevant case according to the data – households hold little or no equity until middle age. Third, the correlation between consumption growth and equity returns rises with age, but it is low at all ages for reasonable parameter choices. Fourth, risk aversion estimates based on the standard excess return formulation of the consumption Euler Equation are greatly upward biased. The bias diminishes, but remains large, for "samples" of households with positive equity holdings. In all four respects, the introduction of a wedge between borrowing costs and the risk-free return greatly improves the fit between theory and empirical evidence.

Table 1 reports data on the size of the wedge. The bottom two rows show that household borrowing costs on unsecured loans exceed the risk-free return by about six to nine percentage points on an annual basis, after adjusting for tax considerations and charge-offs for uncollected loan obligations. Since 1987, roughly two percentage points arise from the asymmetric income tax treatment of household interest receipts and payments. However, the bulk of the wedge arises from transactions costs in the loan market. Despite the evident size of these transaction costs, they have been largely ignored in theoretical analyses of life-cycle consumption and portfolio behavior. They have also been ignored in most empirical studies of asset-pricing behavior.

Aside from the wedge between borrowing costs and the risk-free return, our life-cycle model is entirely standard. None of our results rely on strong risk aversion, a high degree of impatience, habit formation, other types of nonseparability, self-control problems or myopia. Nor do they rely on equity market participation costs, informational barriers, time-varying asset returns, enforcement problems in loan markets, or hard borrowing limits (other than the lifetime budget constraint).

According to our analysis, the cost of borrowing and the shape of the expected life-cycle income profile are key determinants of equity accumulation. The variance of undiversifiable labor income risk also has important effects on equity holdings, and the risky nature of labor income is an important source of cross-sectional heterogeneity in consumption and asset holdings conditional on age.

Our analysis also highlights other interesting comparative static results. First, equity holdings and market participation rates are non-monotonic in the cost of borrowing, both reaching a minimum where the borrowing cost equals the expected return on equity. As this result suggests, the properties of our model are not "in between" those of the canonical model with equal borrowing and lending rates and alternative models with hard borrowing limits. Second, equity holdings and participation rates are non-monotonic in the degree of risk aversion when borrowing costs exceed the risk-free return. Third, once we consider realistic borrowing costs, neither sizable deviations from optimal portfolio shares nor delayed participation in equity markets until, say, age 50 involve large welfare costs.

On the asset-pricing front, our analysis helps resolve two closely related aspects of the equity premium puzzle: Why do people hold so little equity when faced with a substantial equity premium? And, why is the covariance between consumption growth and asset returns so low? Answers to both question follow from the high borrowing costs reported in Table 1. These data imply a small or negative "leverage" premium on borrowed funds that are invested in the stock market. This fact, when combined with upward sloping income profiles or mildly impatient consumers, means that households accumulate little or no equity until middle age, and they can have modest equity holdings even on the verge of retirement.

The paper proceeds as follows. The balance of the introduction discusses related research and reviews some well-established facts about life-cycle consumption and portfolio behavior. Sections 2 and 3 describe the model and the choice of parameters. Section 4 discusses life-cycle portfolio choice and consumption in our model, and section 5 compares implications of the model with empirical evidence. Section 6 concludes by briefly considering some additional implications of the model and related empirical evidence. An appendix describes our numerical solution method.

1.1 Relationship to the Literature

Our model departs modestly from the seminal work on life-cycle portfolio behavior by Merton (1969) and Samuelson (1969). Indeed, our model is identical to Samuelson's discrete-time setup except for three elements: the wedge between borrowing costs and risk-free returns, the presence of undiversifiable income shocks, and the calibration to realistic income profiles. The wedge and the undiversifiable labor income shocks necessitate a computational approach to the analysis, which we pursue using methods developed by Judd, Kubler and Schmedders (2002).

We also build on other research in finance and macroeconomics. Brennan (1971) shows that a wedge between borrowing costs and risk-free returns is easily handled in the standard one-period model of mean-variance portfolio choice. The wedge implies that households cannot attain points above and to the right of the tangency portfolio along the standard capital market line. Higher borrowing costs reduce the demand for equity in the one-period setting, given standard mean-variance preferences. Heaton and Lucas (1997) show that a borrowing rate that exceeds the risk-free return reduces equity holdings for an infinitely lived agent.

We incorporate key life-cycle elements into a many-period setting with realistic borrowing costs and endogenous wealth accumulation. In our model, unlike Brennan or Heaton-Lucas, higher borrowing costs raise the demand for equity in reasonable circumstances. The causal mechanism behind this result involves the impact of borrowing costs on precautionary savings and life-cycle asset accumulation. More generally, life-cycle factors play a central role in both equity market participation and equity accumulation behavior in our model.

Many researchers have explored the effects of hard borrowing limits on portfolio choice in life-cycle models. For example, Constantinides, Donaldson and Mehra (2002) consider a three-period model with no borrowing, and Gomes and Michaelides (2002) consider a calibrated many-period life-cycle model. We show that allowing households to borrow at a rate above the risk-free return yields more realistic behavior with respect to asset accumulation, equity market participation and borrowing itself than a blanket prohibition on borrowing.

Section 5 shows that our model generates a standard deviation of marginal utility growth that is an order of magnitude smaller than the Sharpe ratio. This result appears inconsistent with Hansen and Jagannathan (1991), who argue that the Sharpe ratio represents a *lower* bound on the standard deviation of marginal utility growth. Our result should not come as a complete surprise, however. Other researchers (He and Modest, 1995, and Luttmer, 1996) show that the existence of trading frictions can *potentially* reduce the implied lower bound on the volatility of marginal utility growth. We show that one particular, and easily quantified, friction does indeed reconcile a high Sharpe ratio with low volatility of marginal utility growth.

1.2 Some facts of life-cycle consumption and portfolio choice

Three well-established sets of empirical results are relevant to an assessment of our model. First, a large percentage of households hold no equity – a phenomenon sometimes referred to as the "participation puzzle." According to Vissing-Jorgenson (2002), only 44 percent of households held stock in 1994, a big increase over the 28 percent figure for 1984. Participation rates rise with age (Poterba and Samwick, 2001), education and income (Mankiw and Zeldes, 1991, Brav and Geczy, 1995), and self-employed workers are more likely to hold stock (Heaton and Lucas, 2000). To a large degree, low equity market participation can be traced to the fact the many households have little or no financial wealth (Lusardi et al., 2001).

Second, most households that do hold equity hold very little. Vissing-Jorgensen (2002) reports that the median level of equity holdings for stockholding households is about 21 thousand dollars, and the mean is 95 thousand dollars. Ameriks and Zeldes (2001) find that the level of stockholding rises with education, income and age.

Third, at least as far back as Grossman and Shiller (1982), many researchers have observed that the covariance of consumption growth and equity returns is very low, even for most households that hold equity. That is, given standard preferences and a plausible degree of risk aversion, the estimated covariance violates standard consumption-based asset pricing relations. Several researchers have also shown, however, that the covariance is higher for households that hold equity. See, for example, Mankiw and Zeldes (1991), Brav and Geczy (1995), Brav, Constantinides and Geczy (2002) and Attanasio, Banks and Tanner (2002).

2 A Life-Cycle Model

We consider a partial equilibrium model of household consumption and portfolio choice. The household life cycle consists of two phases, work and retirement, which differ with respect to the character of labor income. During the working years, log labor income ($\tilde{y}_t$) evolves as the sum of a deterministic component (d_t), a random walk component ($\tilde{\eta}_t$), and an uncorrelated transitory shock ($\tilde{\varepsilon}_t$):

$$\tilde{y}_t = d_t + \tilde{\eta}_t + \tilde{\varepsilon}_t.$$

This type of income process is widely used in life-cycle studies of consumption and asset accumulation.

During the retirement years, a household receives a fraction of its income in the last year of work. Ideally, we would specify retirement income as some fraction of, say, the highest n years of work income – consistent with most defined benefit pension plans and with social security. However, such a structure is computationally burdensome, because it increases the dimensionality of the state space. As a computationally easier alternative, we first calculate the ratio of the average value of d_t in the highest n working years to the value of d in the last year of work. We then multiply this ratio by realized income in the last year of work to get the retirement basis. Finally, to get retirement income, we multiply the retirement basis by a number between zero and one called the replacement rate.

To see why this procedure is useful, look at Figure 1, which shows expected labor income profiles for various education groups. Note, for example, that expected labor income in the final year of work is lower relative to lifetime expected income for college-educated households than for households with post-graduate education. Hence, setting retirement income to the same fraction of income in the last working year for all households would effectively assign a relatively low replacement rate to college-educated households. The retirement basis serves to adjusts for differences among education groups in the shape of the lifetime expected income profile.

Households can trade three financial assets. They can buy equity with stochastic net return $\tilde{r}_E$, save at a net risk-free rate r_L, and borrow at the rate $r_B \geq r_L$. Households cannot take short positions in equity, nor can they borrow negative amounts. Net indebtedness cannot exceed the present value of the household's lowest possible future income stream, discounted at the borrowing rate of interest.

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

$$U(C_t) + \mathrm{E}_t \sum_{a=t+1}^{T} \beta^{a-t} U(\tilde{C}_a)$$

subject to a sequence of budget constraints, where C_a is consumption at age a, E_t is the expectations operator conditional on time-t information, β is a time discount factor, and $U(\cdot)$ is an isoelastic utility function. We solve numerically for the optimal solution using a backward induction algorithm, as described in the appendix.

3 Parameter Settings and Discretization

Tables 2 and 3 summarize our parameter settings. Following Campbell (1999), we set the annual risk-free investment return to 2%, the expected return on equity to 8% and the standard deviation of equity returns to 15%. We set the correlation of equity returns and labor income shocks to zero.[1] In line with Table 1, we set the baseline borrowing rate to 8%.

For the life-cycle income processes, we adopt parameter values estimated by Gourinchas and Parker (2002) from the Consumer Expenditure Survey (CEX) and the

[1]Davis and Willen (2000) present evidence of non-zero correlations between labor income shocks and equity returns. They also consider the implications of a non-zero correlation for life-cycle portfolio choice in a model with the borrowing rate equal to the risk-free investment return.

Panel Study of Income Dynamics (PSID). The GP income measure is "after-tax family income less social security tax payments, pension contributions, after-tax asset and interest income" in 1987 dollars. GP also subtract "education, medical care and mortgage interest payments" from their measure of income, because "these categories of expenditure do not provide current utility but rather are either illiquid investments or negative income shocks."[2] They restrict their sample to male-headed households and attribute the head's age and education to the entire household.

To estimate the deterministic component of income for each education group, GP fit a fifth-order polynomial in the head's age to CEX data on log family income. They also fit a fifth-order polynomial to their entire sample, pooled over the five education groups. To estimate the standard deviation of transitory and permanent income shocks, GP use the longitudinal aspect of the PSID. Since the income measures reported in household surveys contain much measurement error, the raw variance estimates substantially overstate income uncertainty. To adjust for this overstatement, we adopt GP's suggestion to reduce the estimated variance of the transitory shock by one half and the variance of the permanent shock by one third. Table 3 reports the standard deviation of the income shocks after adjusting for measurement error.

The expected income profiles displayed in Figure 1 reflect three elements of the GP income processes: (i) the profile of the deterministic component; (ii) the variance of the transitory shock to log income, which affects the level of expected income; and (iii) the variance of the permanent shock, which affects the level and slope of expected income. The profile for the middle education group, not displayed in Figure 1, is very similar to the profile for the pooled sample.

We discretize the state-space using the Tauchen and Hussey (1991) method. Our model has three sources of randomness: a permanent labor income shock, a transitory income shock and an asset return shock. We specify two discrete points for the permanent shock, two points for the transitory shock and three points for the asset return shock, so that the random shocks obey a twelve-state Markov chain.

Our discretization procedure does not generate zero income in any state of nature. In this respect, our specification differs from that of GP. The difference is not innocuous: by assuming that agents have non-zero probability of zero income, GP preclude borrowing.[3] In our setup, households can and do borrow.

[2]Without these deductions, household income would be about 27% higher.

[3]If zero income is possible in the last period of life, households that borrow in the penultimate period run the risk of negative consumption in the final period. With isoelastic utility and relative risk aversion of at least 1, the possibility of negative consumption, no matter how remote, leads to infinitely negative utility. Thus no households borrow in the penultimate period. The same argument extends to earlier periods of life by induction.

In our numerical analysis below, we often turn off one or both labor income shocks. We do this for two reasons: first, to help understand the impact of income uncertainty on equity demand and, second, to show that many of our results do not rest on income uncertainty. Whenever we shut off one or both income shocks, we also readjust the deterministic income component to preserve the same expected income profile.

4 The demand for equity over the life cycle

Before analyzing portfolio choice, we first introduce some terminology and provide intuition. We then explore how four parameters of the household decision problem affect the demand for equity over the life cycle: (1) the borrowing rate, (2) risk aversion, (3) undiversifiable labor income shocks, and (4) the shape of the income profile. Lastly, we turn to the issue of non-participation in equity markets.

Borrowing capacity is the present value of future labor income (including retirement income), when discounted at the borrowing rate, along the lowest possible future income path. The *equity premium* is the difference between the expected return on equity and the risk-free investment return. The *leverage premium* is the difference between the expected equity return and the borrowing rate. When the cost of borrowing exceeds the risk-free investment return, the equity premium exceeds the leverage premium. Hence, the net return on equity depends on the source of funds invested, as depicted in the following table.

Source of funds	Opportunity cost	Net equity return
Liquid wealth	Risk-free return	Equity premium
Borrowing capacity	Borrowing rate	Leverage premium

Given a realistic equity premium (say 6%), a household invests all or much of its liquid wealth in equity. When the leverage premium is positive, a household may also borrow to finance equity holdings, but it holds less equity than an otherwise similar household with greater financial wealth. When the leverage premium is negative, households do not draw on borrowing capacity to invest in equity. That is, with a zero or negative leverage premium, equity demand varies closely (often one-for-one) with liquid wealth. In turn, liquid wealth depends principally on the shape of the lifetime income profile and the strength of consumption-smoothing motives. For example, a household with a sharply upward-sloping income profile saves little, and thus acquires little equity, early in life. The upshot is that the evolution of liquid wealth over the life cycle is a key determinant of equity demand.

4.1 Effect of the borrowing rate

How does the borrowing rate affect the demand for equity over the life cycle? First, a higher borrowing cost lowers borrowing capacity by reducing the present value of labor income. Second, a higher borrowing rate lowers the leverage premium. And third, the borrowing rate affects the evolution of wealth over the life cycle. A low borrowing rate depresses liquid wealth by encouraging greater borrowing for consumption smoothing purposes and by substituting for precautionary wealth holdings that households would otherwise accumulate to smooth transitory income shocks. But a low borrowing rate can also increase liquid wealth: if the leverage premium is positive, borrowing to invest in equity enables the household to increase wealth over time.

As these remarks suggest, there is a non-monotonic relationship between the cost of borrowing and the demand for equity. For example, consider our baseline specification with no labor income risk for a household from the pooled sample. Figure 2 shows its life-cycle equity holdings (averaged over many draws) for alternative borrowing rates. When the borrowing rate equals the risk-free return of 2%, households invest large amounts in equity throughout the life cycle, a result that is not sensitive to the shape of the income profile. Thus, the standard model with $r_B = r_L$ implies equity holdings that greatly exceed what we see in the data.

A borrowing rate of 5% yields much lower equity holdings throughout the life cycle. Why? An increase in the borrowing rate from 2% to 5% implies a reduction in the leverage premium from 6% to 3% and a decline in borrowing capacity. The effect on a very young household is easily understood: since it has no liquid wealth, a smaller leverage premium and lower borrowing capacity mean lower equity demand. Less obviously, the disparity in equity holdings persists into retirement. Two forces are at work. First, households with a non-zero replacement rate still have borrowing capacity in retirement. As shown in Figure 3, households with a positive leverage premium continue to borrow until the year before death.[4] So even in retirement, the size of the leverage premium affects equity demand. Second, a higher leverage premium earlier in life leads, in expectation, to higher wealth accumulation by retirement, as illustrated in Figure 4. A household with a 2% borrowing rate has much greater liquid wealth at retirement than a household with a 5% borrowing rate.

Equity demand behaves differently when the leverage premium is zero or negative. With a borrowing rate equal to 8%, the household from the pooled sample does not invest in equity until age 54 (Figure 2). That is, it never pays to invest borrowing capacity in equity when the leverage premium is zero or negative. However, the

[4] Uncertain longevity coupled with imperfect annuity markets would attenuate or eliminate the incentive to borrow late in life.

household still borrows when young for consumption-smoothing purposes (Figure 3), provided that the cost of borrowing is not too high, so that liquid wealth is negative during much of the life cycle. Although the household from the pooled sample with a zero leverage premium stops borrowing after age 53, asset accumulation and equity demand are smaller than cases with a positive leverage premium.

At borrowing rates above 8%, households hold more equity than at a borrowing rate equal to 8% (Figure 2). Higher borrowing rates lead to less borrowing early in life (Figure 3) and greater liquid wealth at retirement (Figure 4).

Figure 5 shows average demand for equity as a function of the borrowing rate, computed using the age and education weights in the Gourinchas-Parker CEX sample. As seen in Figure 5, average equity demand is smallest when the borrowing rate equals the expected return on equity. At this trough, equity demand in the model with risky labor income is less than 10% of its value in an otherwise identical model with borrowing rate equal to the risk-free return. For the model with no labor income risk, average equity demand at $r_B = \mathrm{E}(\tilde{r}_E)$ is less than 1% of its value at $r_B = r_L$.

Further increases in the cost of borrowing above $\mathrm{E}(\tilde{r}_E)$ *raise* average equity demand. In particular, higher borrowing costs discourage consumption smoothing through the loan market, so that households begin accumulating wealth earlier in the life cycle. Note, however, that the equity demand function is rather flat to the right of $r_B = \mathrm{E}(\tilde{r}_E)$. Hence, when borrowing rates equal or exceed the return on equity, the demand for equity is an order of magnitude smaller than in a standard model with equal borrowing and lending rates.

To sum up, we emphasize three points. First, even a modest wedge between borrowing and lending rates sharply reduces the demand for equity. Second, a borrowing rate equal to the return on equity minimizes the demand for equity. This result is particularly notable since the borrowing rates reported in Table 1 lie near estimates of the expected return on equity. Third, introducing a realistic borrowing rate into an otherwise standard life-cycle model dramatically shrinks the demand for equity.

4.2 Effect of undiversifiable labor income risk

How does undiversifiable labor income risk affect the demand for equity over the life cycle? First, greater income risk makes households with proper preferences effectively more risk averse, which reduces equity demand at given levels of liquid wealth and borrowing capacity. Second, greater income risk intensifies the precautionary saving motive, which encourages wealth accumulation for consumption-smoothing purposes. These two effects work in opposite directions.

Figure 6 shows that the first effect dominates when $r_B = r_L$, so that labor income uncertainty lowers equity holdings. In contrast, the second effect dominates when $r_B = \mathrm{E}(\tilde{r}_E)$. This case differs from the case with $r_B = r_L$ for two reasons. First, when $r_B = \mathrm{E}(\tilde{r}_E)$, younger households hold no equity in the absence of income uncertainty. Hence, they cannot offload risk by reducing their equity holdings, and the first effect is shut off. Second, it is more costly to rely on borrowing as a consumption-smoothing device at a high interest rate, so that the precautionary motive for asset accumulation becomes stronger. As a result, income uncertainty increases equity demand when $r_B = \mathrm{E}(\tilde{r}_E)$.[5]

4.3 Effect of relative risk aversion

How does an increase in the relative risk aversion parameter affect the demand for equity over the life cycle? Greater risk aversion lowers a household's appetite for risk, and its demand for equity, at a given level of liquid wealth. But risk aversion also has a powerful effect on the evolution of liquid wealth over the life cycle. Higher risk aversion means higher precautionary savings, which raises wealth. Higher risk aversion also means a lower elasticity of substitution under our preference specification, which leads to more borrowing and less wealth accumulation with a rising income profile.

As these remarks suggest, stronger risk aversion can mean higher or lower equity demand, and the effects can vary significantly with age. When the borrowing rate equals the risk-free return, higher risk aversion leads to lower equity holdings throughout the life cycle. When the borrowing rate equals the return on equity, the story is more complicated, as shown in Figure 7. Observe that a household with RRA=1 holds less equity throughout life than a household with RRA=.5 or one with RRA=4. A household with RRA=8 holds more equity early in life than any other case shown, but it holds less equity at the end of life.

What drives these results? Since the leverage premium is zero, the main issue is how liquid assets evolve over the life cycle. A household with RRA=0.5 has a high elasticity of substitution, which makes it willing to reduce consumption early in life. As a result, it borrows less early in life than a household with RRA=1. In turn, lower borrowing and more investment lead to higher wealth accumulation throughout the life cycle. A household with RRA=4 accumulates assets early in life because of a strong precautionary motive. The additional savings early in life lead to greater wealth throughout life than the household with RRA=1. Finally, what about the household with RRA=8? Early in life, it invests more than the other cases shown, reflecting its strong precautionary saving motive. But all along, unlike the other

[5] We discuss the distinct effects of permanent and transitory income shocks on the demand for equity in Sections 5.1 and 5.2 below.

households, it finds the risk-free asset highly attractive. As it ages, it continues to direct a large fraction of its portfolio to the risk-free asset, so that it experiences a substantially lower return on its asset portfolio.

Figure 8 shows that the average demand for equity is a non-monotonic function of the risk aversion parameter when borrowing costs exceed the risk-free return. For relative risk aversion below 2 or above 7, equity demand rises with risk aversion, as predicted by a simpler model with $r_B = r_L$. For relative risk aversion between 2 and 7, equity demand is a declining function of the risk aversion parameter. Equity demand is near its minimum for relative risk aversion values near 2 or 3.

4.4 Effect of the expected labor income profile

How does the shape of the expected income profile affect the demand for equity? The answer hinges on the cost of borrowing. When $r_B = r_L$ and labor income is risky (certain), the shape of the income profile has little (zero) effect on equity demand. In contrast, when $r_B \geq \mathrm{E}(\tilde{r}_E)$, the demand for equity is highly sensitive to the shape of the income profile.[6] The explanation for this sensitivity is straightforward: households borrow only for consumption-smoothing purposes when $r_B \geq \mathrm{E}(\tilde{r}_E)$, so they hold no equity until, and insofar as, they maintain positive financial wealth. The age at which this occurs depends on the slope of the income profile.

Consider the case with $r_B = \mathrm{E}(\tilde{r}_E)$. Figure 9 compares life-cycle equity demand for a household from the pooled sample and no labor income risk to an otherwise identical household with a flat income profile. We set income in the flat profile to the simple mean of baseline labor earnings during the working years. The household with a flat profile invests in equity throughout life, whereas the household with the upward-sloping baseline profile waits until age 54. Early investment, compounded by the high return on equity, means that the household with a flat profile accumulates large wealth and equity positions before the baseline household even begins to invest.

Figure 10 shows the effect of the income replacement rate during retirement on equity holdings. In this figure, we vary the replacement rate while holding fixed the other income parameters. The more income a household expects to receive in retirement, the less it saves and the less it invests in equity. The nature of retirement income also affect the optimal portfolio mix. When replacement rates are low, or when pensions take the form of defined contribution plans invested heavily in equities, households are more likely to invest in the risk-free asset in addition to equity.

[6] The shape of the income profile also affects equity demand in the intermediate case with $r_B \in (r_L, \mathrm{E}(\tilde{r}_E))$, but the effect is stronger when $r_B \geq \mathrm{E}(\tilde{r}_E)$.

4.5 Non-participation in equity markets

When the leverage premium is zero or negative, a household invests in equity if and only if its net liquid wealth position is positive. Since households can borrow in our model, net liquid wealth is often negative or zero, and the household holds no equity. This pattern of borrowing and non-participation in equity markets is fully consistent with rational, time-consistent behavior by patient, mildly risk averse households.

Table 4 shows participation rates for different age groups using the pooled-sample income process and baseline parameter settings. When the cost of borrowing equals the return on equity, and labor income is not too risky, equity market participation rates are low or zero during much of life. A borrowing rate above the return on equity *raises* participation for precisely the same reason that it leads to higher asset holdings. As we discuss shortly, the impact of income uncertainty on equity market participation depends on the persistence of the income shocks.

Not shown in the table is the ambiguous effect of risk aversion on equity market participation. Participation rates are high for very low levels of risk aversion ($RRA < 1$) and for high levels ($RRA > 4$), but they are considerably lower for intermediate levels ($1 \leq RRA \leq 4$). The explanation for the non-monotonic relationship between participation and risk aversion parallels the explanation given above for the non-monotonicity in the level of equity holdings.

5 Does our model fit the facts?

Section 1 reviewed several facts about life-cycle consumption and portfolio behavior. We now assess the fit between those facts and predictions of the model. Where the model fails to fit the facts, we assess whether the failure is large or small in a welfare sense.

5.1 Fact 1: Non-participation in equity markets

To get a better sense of what drives participation behavior in the model, Table 5 shows how average participation rates vary with several aspects of the specification. Two points are immediately clear from Tables 4 and 5. First, the model can generate low participation rates when borrowing rates exceed the risk-free return. Second, when calibrated to the GP data on income shock variances, the predicted participation levels exceed those observed in the data.

In evaluating the failure of the model with the GP income process to match stylized facts on participation, it bears repeating that we have not modelled any liquidity advantage for the risk-free asset. Given the modest liquid wealth positions of many households in our simulations, equity market participation would be much lower if we incorporated some reason to hold small liquid wealth positions in the risk-free asset. We make a related point below when we quantify the welfare costs of non-participation from the vantage point of the model.

Tables 4 and 5 also show that permanent and transitory income shocks affect participation quite differently. The introduction of permanent income shocks raises participation in all scenarios shown, and specifications with only permanent shocks exhibit full participation at all ages. In contrast, transitory income shocks push outcomes away from zero and 100% participation. Hence, relative to a specification with no income risk, transitory income shocks tend to raise participation at younger ages. But relative to a specification with permanent income shocks, the introduction of transitory income shocks lowers participation at younger ages.

The explanation for these results involves the consumption-smoothing role of borrowing. Borrowing is not helpful in smoothing permanent income shocks, but it is useful for smoothing transitory shocks. In particular, a sufficiently bad transitory shock (or shock sequence) causes the household to draw down liquid wealth and resort to borrowing, at which point it ceases to hold equity. The consumption smoothing role of borrowing is highly sensitive to the the borrowing rate and the higher the borrowing rate, the more households households act to reduce the likelihood of borrowing. This behavioral response can be seen in Table 4 as a positive relationship between r_B and participation, conditional on age, when only transitory income shocks are present. In short, transitory shocks create a motive to hold equity when the household would otherwise hold none, but they also give rise to circumstances in which some households exhaust their asset holdings and turn to borrowing.

Turning to cross-sectional evidence, empirical studies consistently find that equity market participation rates rise with age.[7] As seen in Table 4 and several of the figures, this empirical regularity is well matched by the predicted life-cycle pattern of participation. In this respect, our analysis provides a simple explanation for a widely observed empirical regularity.

Empirical studies also find that participation rates rise with education. In contrast, our specifications predict lower participation rates for the top two education groups, because they face steeper expected income profiles. It is unclear whether this mismatch reflects a failure of the model or an inadequate calibration. For reasons of

[7] Isolating age effects from time and cohort effects requires an identifying assumption. However, to the best of our knowledge, every study that considers the issue concludes that age has a positive effect on participation in equity markets.

data availability, our simulations confront all households with the same environment except for the differences among education groups in the income processes. However, small differences among education groups in borrowing costs, risk aversion and time discounting, or small changes in the baseline values for the income shock variances, can alter the predicted relationships between participation rates and education.

5.2 Fact 2: Average and life-cycle demand for equity

For realistic borrowing rates, our model predicts small equity holdings, roughly in line with the data. The impact of borrowing costs on average equity demand can be seen in Table 5. The column headed "DKW" shows predicted average equity holdings when $r_B = \mathrm{E}(\tilde{r}_E)$, while the column headed "Std." shows average equity holdings in an otherwise identical model with $r_B = r_L$. Comparing these two columns shows that the borrowing cost wedge dramatically reduces equity demand in every specification.

Focusing now on cases with $r_B = \mathrm{E}(\tilde{r}_E)$, some specifications imply tiny values for average equity demand, as illustrated by the row with RRA=4 and no labor income risk. Specifications with risky labor income imply substantially larger equity holdings, in line with the discussion in Section 4.2

Table 6 shows that a realistic borrowing rate leads to much lower equity holdings at all ages. Table 6 and Figures 6 and 7 show that equity holdings rise sharply with age until late in the life cycle when $r_b = \mathrm{E}(\tilde{r}_E)$.

5.3 Fact 3: Covariance of consumption and equity returns

How does our model fit the facts about the covariance of consumption growth and equity returns? Much better, in three respects, than models with equal borrowing and lending rates. First, our model's implied covariance between minus marginal utility growth and equity returns lies well below the excess return on equity. Second, if one applies standard formulas to calculate risk aversion using data generated by our model, one overestimates risk aversion, often by a factor of ten or more. This implication rationalizes the implausibly large estimates of risk aversion in consumption-based asset-pricing studies. Third, our model rationalizes the rejection of standard consumption-based asset-pricing relationships in samples that are restricted to households with positive equity holdings. We now develop these three points in turn.

First, consider the covariance between minus marginal utility growth and equity returns. According to standard models with equal borrowing and lending rates,

$$\mathrm{E}\left(\tilde{r}_E\right) - r_L = \mathrm{cov}\left(\tilde{r}_E, -\frac{\Delta\widetilde{MU}}{\mathrm{E}(\Delta\widetilde{MU})}\right) \equiv \mathrm{cov}_{MU} \tag{1}$$

That is, the excess return on equity equals minus the covariance between equity returns and the growth rate of marginal utility. Equation (1) is central to most consumption-based asset pricing studies. It fares poorly in confrontations with the data for standard preference specifications.

As shown in Tables 5 and 6, equation (1) fails to hold in versions of our model with realistic borrowing costs. Table 6 shows that cov_{MU} is (near) zero for young households and reaches a maximum of 2.3 prior to retirement in the version of the model with risky labor income. In contrast, the standard model implies that cov_{MU} equals 6.0 at all ages.

We can interpret cov_{MU} as the equity premium implied by consumption behavior. In particular, suppose we observe a consumption process and know the utility function. In addition, suppose we assume that households face equal borrowing and lending rates. Then from equation (1), we would infer an equity premium of 0.2% in our baseline specification with no labor income risk and 1.1% with risky labor income (first and third rows of Table 5). In other words, we would mistakenly attribute the low covariance of consumption growth and asset returns to a low equity premium, not differential borrowing and lending rates.

We can also follow Hansen and Jagannathan (1991) and interpret the same information in another way. Equation (1) implies that the standard deviation of marginal utility growth is bounded below by the Sharpe ratio. We can use our model to calculate the actual standard deviation of marginal utility growth, reported in Table 5 in the "HJ" column. Note that only when the borrowing cost wedge is zero or very small does the HJ statistic equal or exceed the Sharpe ratio (38%). For realistic borrowing costs, the model implies that the standard deviation of marginal utility growth is much smaller than the Sharpe ratio.

Second, consider the estimation of risk aversion from consumption-based asset-pricing relationships. Recall from Hansen and Singleton (1983) that, if asset returns and consumption are jointly log-normal and individual utility is of the time-separable isoelastic form, then equation (1) implies

$$\gamma = \frac{\mathrm{E}(\tilde{R}) - R_f + .5\,\mathrm{var}(\tilde{R})}{\mathrm{cov}(\tilde{R}, \tilde{c})} \tag{2}$$

where γ is the coefficient of relative risk aversion. The columns headed "Implied RRA" in Tables 5 and 6 calculate risk aversion using equation (2).

Inspecting these tables reveals that equation (2) greatly overestimates risk aversion when borrowing rates exceed lending rates. In the baseline specification with permanent and transitory income shocks, the Implied RRA is 13.0, more than six times the true value. In the baseline specification with certain labor income, the Implied RRAA is 60.9, more than thirty times its true value.

Third, consider the consequences of restricting samples to households with positive equity holdings. Starting with Mankiw and Zeldes (1991), researchers have argued that equations (1) and (2) are unreasonable descriptions of reality for households that do not participate in equity markets. But, as the argument goes, perhaps equations (1) and (2) approximate reality for households that invest non-zero amounts in equity. This type of argument leads to equity ownership as a sample selection criterion in studies of consumption-based asset pricing relationships.

To evaluate this approach in terms of our model, the bottom row of Table 6 reports statistics for "samples" restricted to equity market participants. This selection criterion leads to a better fit between theory and data, but the theory continues to perform poorly if we maintain that the data are generated by a world with $r_B = r_L$. For example, restricting the sample to equity market participants lowers the Implied RRA from 55.8 to 15.7 in the specification with no income risk, as compared to a true RRA value of 2.0. If one instead recognizes that the borrowing rate is several percentage points above the risk-free return, then the theory explains why consumers appear so risk averse in consumption-based asset pricing studies, why the implicit equity premium appears so small, and why the Hansen-Jagannathan bound is violated.

5.4 Facts we don't match

Our model fails to match the facts on at least two dimensions. First, as we discussed above, versions of the model with risky labor income predict higher equity market participation rates than seen in the data. Second, the model predicts that many households hold equity but not the risk-free asset. Moreover, unless risk aversion is high, the model predicts small portfolio shares in the risk-free asset for households that invest in both assets.

We now investigate the welfare significance of these model failures. Our basic message is simple: When borrowing rates are in the neighborhood of the return on equity, the welfare costs of delayed participation in equity markets or the wrong bond-equity mix are small. Even the right to hold equity can have modest welfare consequences in a life-cycle model with a zero or negative leverage premium. The situation is very different in the standard model. When the borrowing rate equals the risk-free investment return, delayed equity market participation and the wrong bond-equity mix imply large welfare costs. These observations are not intended to downplay the shortcomings of our model. Rather, they suggest that there is ample scope for other factors – participation and transaction costs, alternative preferences, a desire for liquidity, information costs and so on – to strongly influence equity holdings, once we recognize that borrowing costs exceed the risk-free return.

To carry out the welfare comparisons, it is useful to focus on certainty-equivalent consumption levels, which we report in Table 7. To obtain certainty-equivalent consumption, we first calculate lifetime expected utility, U, for a given consumption profile. We then find the constant level of consumption, $\overline{c}$, that yields the same level of lifetime expected utility. That is, we solve

$$\sum_{t=0}^{T} \beta^t \frac{\overline{c}^{1-\gamma}}{1-\gamma} = U \quad \text{for} \quad \overline{c} = \big[\frac{1-\gamma}{\sum_{t=0}^{T} \beta^t} U\big]^{\gamma-1},$$

where β is the time discount factor, and γ is relative risk aversion.

For example, consider a household with less than high school education, no income risk and an 8% borrowing rate. According to Table 7, that household is indifferent between the optimal lifetime consumption plan and a flat profile with consumption of \$17,850 per year. If the same household is not allowed to trade equity, certainty-equivalent consumption falls to \$17,800. In other words, the right to trade equity is worth about \$50 per year, or less than .3 percent of annual consumption. Table 7 considers three departures from optimal behavior: non-participation in equity markets throughout life, non-participation until age 50 with optimal behavior thereafter, and optimal participation subject to a 50-50 bond-equity portfolio mix.

Perhaps the most notable feature of Table 7 is the small cost of delayed equity market participation, given realistic borrowing costs. In the specification with no income risk and an 8% borrowing rate, the cost of waiting until age 50 to invest in equity is essentially zero. The costs of delayed participation are larger for specifications with risky income, but less than \$400 per year for every group. Costs on this order are not trivial, but they seem well within range of reasonable values for the costs of participating and transacting in equity markets.

The standard life-cycle model with $r_B = r_L$ tells a starkly different story. For example, the costs of delayed participation until age 50 when $r_B = r_L$ range from 14 to 23 thousand dollars per year, which amounts to more than fifty percent of annual consumption. In other words, the standard model implies enormous welfare costs for widely observed departures from theoretically predicted behavior. These costs are far too large to be rationalized by plausible costs of participating and transacting in equity markets.

6 Conclusion

Realistic borrowing costs dramatically shrink the demand for equity in the canonical life-cycle model of Merton and Samuelson. Given realistic borrowing costs, theoretical implications about the average level and life-cycle pattern of equity holdings are roughly in line with the data. In addition, a realistic borrowing rate greatly improves the performance of the theory along several dimensions that feature prominently in consumption-based asset pricing studies.

The implications of our model differ in noteworthy respects from those in models that abstract from life-cycle considerations and models with hard borrowing limits. For example, equity demand and equity market participation rates are non-monotonic in borrowing costs and risk aversion in our model. The opportunity to borrow at realistic rates in a life-cycle setting also has important consequences for life-cycle wealth accumulation, which in turn drives the demand for equity. Younger households facing an upward sloping income profile borrow against future income, which delays the date at which they first invest in equity or first accumulate substantial equity holdings.

Our analysis points to several directions for future research. We mention three here. First, it provides a new perspective on life-cycle wealth accumulation. Our analysis implies that most households accumulate little or no liquid wealth until middle age, which is consistent with much empirical evidence (e.g., Lusardi, Cossa and Krupka, 2001). Given its simplicity and its assumption of patient, time-consistent, rational consumers, our model and analysis challenge claims that households save too little. Second, our analysis puts the spotlight on the role of borrowing costs and leverage as key factors in the demand for risky assets. This paper does not consider other ways to leverage equity investments such as margin loans or investments in levered mutual funds. Margin loans provide limited scope for levered equity holdings (because of high legal margin requirements), but corporate bonds, government securities, real estate and small business wealth are often subject to less stringent restrictions on leverage. Kubler and Willen (2002) pursue an extended version of our model to address portfolio choice in a broader setting that encompasses a fuller menu of risky assets and leveraging methods. Third, our model has strong implications for the life-cycle profiles of equity holdings and the covariance between consumption growth and equity returns. As a household ages, our model predicts rising equity holdings and an increasing covariance between consumption and equity returns.

References

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

[2] Ameriks, John and Stephen P. Zeldes, 2001, How do household portfolio shares vary with age? working paper, Columbia University.

[3] Attanasio, Orazio P., James Banks and Sarah Tanner, 2002, Asset holding and consumption volatility, *Journal of Political Economy,* 110(4), 771-792.

[4] Brav, Alon, George M. Constantinides and Christopher C. Geczy, 2002, Asset pricing with heterogeneous consumers and limited participation: Empirical evidence, *Journal of Political Economy,* 110(4), 793-824.

[5] Brav, Alon and Christopher C. Geczy, 1995, "An empirical resurrection of the simple consumption CAPM with power utility," working paper, University of Chicago.

[6] Brennan, Michael J, 1971, Capital market equilibrium with divergent borrowing and lending rates, *Journal of Financial and Quantitative Analysis,* 6, 1197-1205.

[7] Campbell, John Y, 1999, Asset prices, consumption and the business cycle, in John B. Taylor and Michael Woodford, editors, *Handbook of Macroeconomics, Volume 1C.* Amsterdam: North-Holland.

[8] Constantinides, George M., John B. Donaldson and Rajnish Mehra, 2002, Junior can't borrow: A new perspective on the equity premium puzzle, *Quarerly Journal of Economics,* 117(1), 269-296.

[9] Davis Steven J. and Paul S. Willen, 2000, Using financial assets to hedge labor income risks, working paper, University of Chicago.

[10] Garcia, C.B. and W.I. Zangwill, 1981, *Pathways to Solutions, Fixed Points, and Equilibria.* Englewood Cliffs, NJ.: Prentice-Hall.

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

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

[13] Grossman Sanford and Robert Shiller, 1982, Consumption correlatedness and risk measurement in economies with non-traded assets and heterogeneous information, *Journal of Financial Economics,* 10(2), 195-210.

[14] Hansen, Lars Peter and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy,* 99(2), 225-62.

[15] Hansen, Lars Peter and Kenneth Singleton, 1983, Stochastic consumption, risk aversion, and the temporal behavior of asset returns, *Journal of Political Economy,* 91, 249-265.

[16] He, Hua and David M. Modest, 1995, Market frictions and consumption-based asset pricing, *Journal of Political Economy,* 103(1), 94-117.

[17] Heaton, John and Deborah Lucas, 1997, Market frictions, savings behavior, and portfolio choice, *Macroeconomic Dynamics* 1, 76-101.

[18] Heaton, John and Deborah Lucas, 2000, Portfolio Choice and asset prices: The importance of entrepreneurial risk, *Journal of Finance,* 55(3), 1163-98.

[19] Judd, Kenneth, Felix Kubler and Karl Schmedders, 2002, Computational methods for dynamic equilibria with heterogeneous agents. Forthcoming in M. Dewatripont, L. P. Hansen and S. J. Turnovsky, editors, *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress.* Cambridge University Press.

[20] Kubler, Felix and Paul Willen, 2002, Margin borrowing and life-cycle portfolio choice. Work-in-progress.

[21] Lusardi, Annamaria, Ricardo Cossa and Erin L. Krupka, 2001, Savings of young parents, *Journal of Human Resources,* 36(4), 762-794.

[22] Luttmer, Erzo G. J., 1996, Asset pricing in economies with frictions, *Econometrica,* 64(6), 1439-67.

[23] Mankiw, N. Gregory and Stephen P. Zeldes, 1991, The consumption of stockholders and nonstockholders, *Journal of Financial Economics,* 29, 97-112.

[24] Merton, Robert C, 1969, Lifetime portfolio selection under uncertainty: The continuous time case, *Review of Economics and Statistics,* 51, 245-257.

[25] Poterba, James M. and Andrew A. Samwick, 2001, Portfolio allocations over the life cycle, in Seiritsu Ogura, Toshiaki Tachibanaki, and A. David Wise, editors, *Aging Issues in the United States and Japan.* Chicago: University of Chicago Press.

[26] Samuelson, Paul A, 1969, Lifetime portfolio selection by dynamic stochastic programming, *Review of Economics and Statistics,* 51, 239-243.

[27] Tauchen, George and Robert Hussey, 1991, Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica,* 59(2), 371-96.

[28] Vissing-Jorgensen, Annette, 2002, Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures. working paper, University of Chicago.

[29] Watson, L.T., S.C. Billups and A.P. Morgan, 1987, HOMPACK: A suite of codes for a globally convergent homotopy algorithm. *ACM Transactions on Mathematical Software,* 13, 281-310.

Appendix: Computational details

We use the standard backward induction to solve the model. At age 80, the solution is trivial: consume everything and invest nothing in any assets; we then solve for optimal consumption and portfolio choice at age 79 conditional on liquid wealth, income and the state of the world and the (degenerate) policy rule at 80; we solve for consumption and portfolio choice at age 78 again conditional on liquid wealth, income, the state of the world and the calculated optimal policy rules at 79. And so on.

The problem therefore reduces to solving simple two-period optimization problems and to approximating policy-rules as functions of some minimal set of states and the current age. In this appendix, we focus on two aspects of the solution: how to reduce the endogenous state space to one variable and how to solve the two period optimization problem effectively. First, we show that we can use preference homotheticity to simplify the problem and reduce the number of continuous state variables to one. Second, we discuss how we solve the two-period optimization problem by solving the first order conditions using homotopy methods.

A little notation will help. Let z_t follow a Markov chain with finite support $z \in \{1, ..., S\}$ and transition π. Gross equity returns equal $\tilde{R}_E(z_t)$, and the gross borrowing and lending rates are R^B and R^L respectively. A date-event z^t is a history of shocks $(z_1, ..., z_t)$. Let $y(z^t)$ equal income at time t.

Preference homotheticity allows us to simplify the problem by combining wealth and income into one variable, an insight due to Deaton (1991). Suppose we have solved for optimal policy rules from time $t+1$ on. Suppose at date t, we are in state z with income y_t and wealth Ξ. Our policy rule for next period specifies that we invest $F^i_{z',t+1}(\Xi_{t+1}(z'), y_{t+1}(z'))$ in asset $i = B, L, E$ at time $t+1$. Bellman's principle implies that the solution to the two-period problem below constitutes optimal portfolio choice at t in state z with income y and liquid wealth Ξ.

$$
\begin{aligned}
\max_{F^L, F^B, F^E} & \frac{c_t^{1-\gamma}}{1-\gamma} + \beta\, \mathrm{E}\big(\frac{c_{t+1}^{1-\gamma}}{1-\gamma}\big) \qquad (3) \\
\text{st } c_t &= y + \Xi - F^L + F^B - F^E \\
c_{t+1}(z') &= y(z', t+1) + \Xi(z') - F^L_{z',t+1}(\Xi(z'), y_{t+1}(z')) + \\
& \quad F^B_{z',t+1}(\Xi(z'), y_{t+1}(z')) - F^E_{z',t+1}(\Xi(z'), y_{t+1}(z'))) \quad \forall z' \in \{1, ..., S\} \\
\Xi(z') &= F^L R_L - F^B R_B + F^E R_E(z') \quad \forall z' \in \{1, ..., S\} \\
F^L &\geq 0, \quad F^B \geq 0, \quad F^E \geq 0
\end{aligned}
$$

A difficulty here is that the F^i's are functions of two variables, liquid wealth and current income.

Note that if we divide the problem through by y_t, we have an almost identical problem. Specifically, let $x_i = \frac{c_i}{y_t}$. Consider the two period optimization problem:

$$\begin{aligned}
&\max_{f^L,f^B,f^E} \frac{x_t^{1-\gamma}}{1-\gamma} + \beta\,\mathrm{E}\big(\frac{x_{t+1}^{1-\gamma}}{1-\gamma}\big) \qquad (4)\\
&\text{st } x_t = \xi - f^L + f^B - f^E\\
&x_{t+1} = \frac{y(z',t+1)}{y_t}\big[\xi(z') - f^L_{z',t+1}(\xi(z')) + f^B_{z',t+1}(\xi(z')) - f^E_{z',t+1}(\xi(z'))\big]\\
&\xi(z') = \frac{y(z',t+1) + f^L R_L - f^B R_B + f^E R_E(z')}{y(z',t+1)}\\
&f^L \geq 0, \quad f^B \geq 0, \quad f^E \geq 0
\end{aligned}$$

Now notice that the policy rules are function of a single endogenous state variable ξ, the ratio of liquid wealth plus current income and current income. This reduction in the state-space dramatically simplifies computation. And we can recover the solution to the original problem (equation (3)) by multiplying the solution to the transformed problem (equation (4)) by current income:

$$\begin{aligned}
c_t &= y_t x_t\\
F_t^L &= y_t f^L, \quad F^B = y_t f^B, \quad F^E = y_t f^E
\end{aligned}$$

We solve the two-period problem by solving a non-linear system of equations which is necessary and sufficient for agents' optimality. Through a simple change of variables (see Garcia and Zangwill (1981)) we can eliminate all inequalities in the Kuhn Tucker conditions and state the optimality conditions as a system consisting solely of equations. The resulting system has 3 unknowns. Let $\eta_j \in \Re$ for $j = 1, 2, 3$ and define the Kuhn Tucker multiplier for asset j $\mu_j = (\max\{0, -\eta_j\})^3$ and the holding $\theta_j = (\max\{0, \eta_j\})^3$. Note that θ and μ are twice continuously differentiable and that the following relations hold:

$$(\max\{0,\eta_j\})^3 = \begin{cases} (\eta_j)^3 & \text{if } \eta_j > 0 \\ 0 & \text{if } \eta_j \leq 0 \end{cases}$$

$$(\max\{0,-\eta_j\})^l = \begin{cases} 0 & \text{if } \eta_j > 0 \\ |\eta_j|^3 & \text{if } \eta_j \leq 0 \end{cases}$$

Moreover, the complementary slackness conditions hold:

$$(\max\{0,\eta_j\})^3 \geq 0, (\max\{0,-\eta_j\})^3 \geq 0, \text{ and } (\max\{0,\eta_j\})^3 \cdot (\max\{0,\eta_j\})^3 = 0.$$

We implement the algorithm using Fortran 90. A simple Newton method usually works well as a non-linear equation solver when a good starting point is known. In some cases we need to use homotopy methods (as implemented in HOMPACK, see Watson et al (1987)) to solve the system.

Finally we draw attention to two pracitcal aspects of our computational solution. First, the range of $f^j_{t,z}(\xi)$ will generally depend on t and z. In practice we vary the domain only with t, assuming that each period agents save at most 30 percent of their average income and that agents never borrow more than they can pay back throughout their lifetime under the assumption that their income is deterministic and that they never consume less than 2/3 of current income. In simulations we then verify that these bounds are never binding. Second, in generating $f^j_{t,z}(\xi)$ we don't solve problem (4) for every possible value of ξ. We solve problem (4) for some finite number of values of ξ and then use cubic spline interpolation to fill in the rest (see Judd et al (2002) for details on spline interpolation). Since the true policy functions have non-differentiabilities we use 50 knots for each spline interpolation to obtain sufficient accuracy.

Maximal relative errors in Euler equations lie below 10^{-6}. Running times clustered around four or five minutes but range from 2 minutes for models with no labor income risk and borrowing rates above the expected return on equity to about 15 minutes for models with labor income risk and borrowing rates below the expected return on equity.

Table 1: Household Borrowing Costs and Risk-Free Returns, Selected Years [tabx6]

	Year	1972	1980	1984	1987	1990	1995	2001
(1)	Average rate on two-year personal loans	12.5	15.5	16.5	14.2	15.5	13.9	13.2
(2)	Average marginal tax subsidy on borrowing	.181	.247	.249	0	0	0	0
(3)	After-tax borrowing cost (1 - row 2)*(row 1)	10.2	11.6	12.4	14.2	15.5	13.9	13.2
(4)	Rate on three-year U.S. Treasury Securities	5.7	11.5	11.9	7.7	8.3	6.3	4.1
(5)	Average marginal tax rate on interest income	.313	.428	.330	.279	.250	.282	.297
(6)	After-tax risk-free return (1 - row 5)*(row 4)	3.9	6.6	8.0	5.5	6.2	4.5	2.9
(7)	Pre-tax wedge between borrowing cost and risk-free return (row 1 - row 4)	6.7	4.0	4.6	6.5	7.2	7.7	9.1
(8)	After-tax wedge between borrowing cost and risk-free return (row 3 - row 6)	6.3	5.1	4.4	8.7	9.3	9.5	10.3
(9)	Charge-off rate on loans, net of recoveries				0.8	1.0	0.7	1.3
(10)	After-tax wedge net of charge offs (row 8 - row 9)				7.9	8.2	8.8	9.1
(11)	After-tax wedge net of charge offs, credit cards				9.0	8.5	7.9	6.5

Sources: Rows (1) through (8) for 1972 to 1987 are reproduced from Table 1 in Altig and Davis (1992). Data for later years as follows: Rows (1) and (4) are from various issues of the *Federal Reserve Bulletin* and the Federal Reserve's *Annual Statistical Digest*. Row (2) reflects the Tax Reform Act of 1986, which eliminated the tax deductibility of interest payments on non-mortgage loans. Row (5) is from Table 1 in Poterba (2001), which is calculated from the NBER TAXSIM model. Poterba's 1999 value is used for the 2001 entry in row (5). Row (9) is from www.federalreserve.gov/releases/chargeoff/chg_all_sa.txt (visited 3 April 2002). Other rows are calculated by the authors as indicated.

Notes: Borrowing costs, returns and charge-offs are expressed as annual percentage rates. Row (9) reports the value of loans removed from the books and discharged against loan loss reserves net of recoveries as a percentage of loans outstanding. Rows (10) and (11) show the difference between the household cost of borrowing and the rate of return on risk-free investments after adjusting for tax considerations and the charge-off rate. Row (11) is calculated in the same manner as row (10), except that it makes use of interest rate and charge-off data for credit cards instead of two-year personal loans.

Table 2: Parameter Settings

Parameter	Baseline	Alternative values
Relative Risk Aversion	2	0.5, 1, 4, 8
Annual Discount factor	0.95	
Age of labor force entry	21	
Age of retirement	65	
Age of death	80	
var $(\Delta\tilde{\eta})$	see Table 3	
cov $(\Delta\tilde{\eta}, \tilde{r}_E)$	0	
var $(\tilde{\varepsilon})$	see Table 3	
cov $(\tilde{\varepsilon}, \tilde{r}_E)$	0	
Replacement rate	80%	20%, 40%, 60%, 100%
r_L	2%	
r_B	8%	2%, 5%, 8%, 20%, 99%
$\mathrm{E}(\tilde{r}_E)$	8%	
std $(\tilde{r}_E)$	15%	

Table 3: Income Shock Standard Deviations, percent per year

Education group	$\sigma(\tilde{\varepsilon})$ (transitory shock)	$\sigma(\Delta\tilde{\eta})$ (permanent shock)
Some high school	23	12
High school diploma	15	14
Some college	13	12
College degree	14	10
Graduate school	16	9
Pooled Sample	15	12

Note: These values are from Gourinchas and Parker (2002), adjusted for measurement error as described in the text.

Table 4: Equity Market Participation Rates [tabx1]

Age Group	No income shocks $r_B =$			Transitory shocks only $r_B =$			Permanent and transitory shocks $r_B =$		
	8	20	99	8	20	99	8	20	99
22 - 24	0	0	0	0	49	65	65	69	69
25 - 29	0	0	0	1	57	80	84	93	94
30 - 34	0	0	0	0	68	85	96	99	99
35 - 39	0	57	0	1	78	88	100	100	100
40 - 44	0	100	100	5	88	93	100	100	100
45 - 49	0	100	100	20	95	97	100	100	100
50 - 54	20	100	100	54	99	99	100	100	100
55 - 59	100	100	100	87	100	100	100	100	100
60 - 79	100	100	100	100	100	100	100	100	100
All ages	27	64	57	32	82	90	94	96	96

Notes:

1. Table entries are for the pooled sample and, unless otherwise noted, the baseline parameter settings. See Tables 2 and 3.
2. The participation rate is 100% at all ages for $r_B \in \{8, 20, 99\}$ in the specification with permanent income shocks only.

Table 5: Average Equity Demand and other Statistics, Alternative Specifications [tabx3]

r_B	RRA	RR	Labor income shocks?	Income Profile	cov_{MU}	HJ	cov_C	Equity Demand DKW	Equity Demand Std	Implied RRA	% Ptcp
-	-	-	No	-	0.2	1.5	0.1	15	1871	60.9	32
-	-	-	T	-	0.3	1.9	0.1	21	737	47.0	39
-	-	-	P+T	-	1.1	7.1	0.5	80	831	13.0	90
2	-	-	No	-	6.0	38.2	2.8	1871	1871	2.5	100
2	-	-	T	-	6.1	38.8	2.7	737	737	2.6	100
2	-	-	P+T	-	6.1	38.7	2.7	831	831	2.6	100
99	-	-	No	-	0.4	2.3	0.2	27	1871	39.3	64
99	-	-	T	-	0.6	3.8	0.3	40	737	23.6	91
99	-	-	P+T	-	1.2	7.5	0.6	84	831	12.3	97
-	0.5	-	No	-	0.6	4.1	1.3	158	1871	5.4	99
-	0.5	-	P+T	-	0.7	4.3	1.4	177	831	5.1	99
-	4	-	No	-	0.1	0.9	0.0	3	1871	206.9	17
-	4	-	P+T	-	2.8	18.1	0.6	99	831	11.1	94
-	8	-	No	-	0.0	0.2	0.0	0	1871	1424.1	10
-	8	-	P+T	-	5.0	31.8	0.6	104	831	11.4	98
-	-	-	No	Flat	0.9	5.8	0.5	77	1920	15.3	100
-	-	0.2	No	-	0.7	4.6	0.4	49	1644	19.5	45
-	-	1	No	-	0.1	0.7	0.1	7	1937	123.6	26

Notes:

1. Table entries are for the pooled sample and, unless otherwise noted, the baseline parameter settings. See Tables 2 and 3. RR denotes the income replacement rate during retirement.
2. $\text{cov}_{MU} = \text{cov}(\tilde{r}_E, -\Delta\widetilde{MU}/\text{E}(\Delta\widetilde{MU}))$. See Section 5.3 for details.
3. $HJ = \text{std}(\Delta\widetilde{MU}/\text{E}(\Delta\widetilde{MU})$. See Section 5.3 for details.
4. $\text{cov}_C = \text{cov}(\Delta C, \tilde{r}_E)$.
5. DKW denotes our model with indicated value of r_B. Std denotes the standard model with $r_B = r_L = 2\%$ and $RRA = 2$. The equity demand values are in 1987 dollars.
6. The Implied RRA is calculated using equation (2). See Section 5.3.

Table 6: Equity Demand and other Statistics over the Life Cycle [tabx7]

Age Group	*No labor income shocks*				*Permanent and transitory shocks*			
	cov_{MU}	Equity Demand		Implied	cov_{MU}	Equity Demand		Implied
		DKW	Standard	RRA		DKW	Standard	RRA
23 - 29	0.0	0	1263	Inf	0.1	8	240	147.6
30 - 39	0.0	0	1619	Inf	0.6	32	495	16.1
40 - 49	0.0	0	2117	Inf	1.2	81	884	8.7
50 - 59	0.1	9	2475	116.4	1.9	154	1320	5.6
60 - 69	0.8	60	2500	12.5	2.3	215	1633	4.5
70 - 79	1.1	47	1668	8.7	2.2	138	1138	4.5
All ages	0.2	11	1891	55.8	1.1	84	828	9.1
Participants Only	0.7	41	1891	15.7	1.2	88	828	9.1

Note: See notes to Table 5.

Table 7: The Welfare Costs of Sub-Optimal Equity Holdings. [`tabx11`]

Certainty-equivalent consumption levels, thousands of 1987 dollars.

r_B	Income Shocks	Equity Holdings	Less than High School	High School Degree	Some College	College Degree	Graduate School
8	No	No	17.80	23.43	25.30	28.06	26.91
		50/50	17.81	23.44	25.31	28.08	26.93
		Age 50	17.85	23.49	25.37	28.17	26.97
		Optimal	17.85	23.49	25.37	28.17	26.97
2	No	No	21.17	28.34	30.55	33.27	33.59
		Age 50	23.22	31.09	33.52	36.50	36.84
		Optimal	37.70	50.48	54.42	59.28	59.84
8	Yes, both	No	13.32	16.04	18.20	22.83	22.46
		50/50	13.65	16.63	18.68	23.11	22.64
		Age 50	13.58	16.43	18.59	23.17	22.73
		Optimal	14.14	17.35	19.38	23.63	22.93

Notes:

1. Unless otherwise noted, the table entries are for the baseline parameter settings. See Tables 2 and 3.

2. "50/50" indicates a portfolio mix that is constrained to have 50% of its value in equity and 50% in bonds. "Age 50" means that the household is constrained from participating in the equity market until age 50. See Section 5.4 for details.

Figure 1: Expected labor income profiles. `figa7`

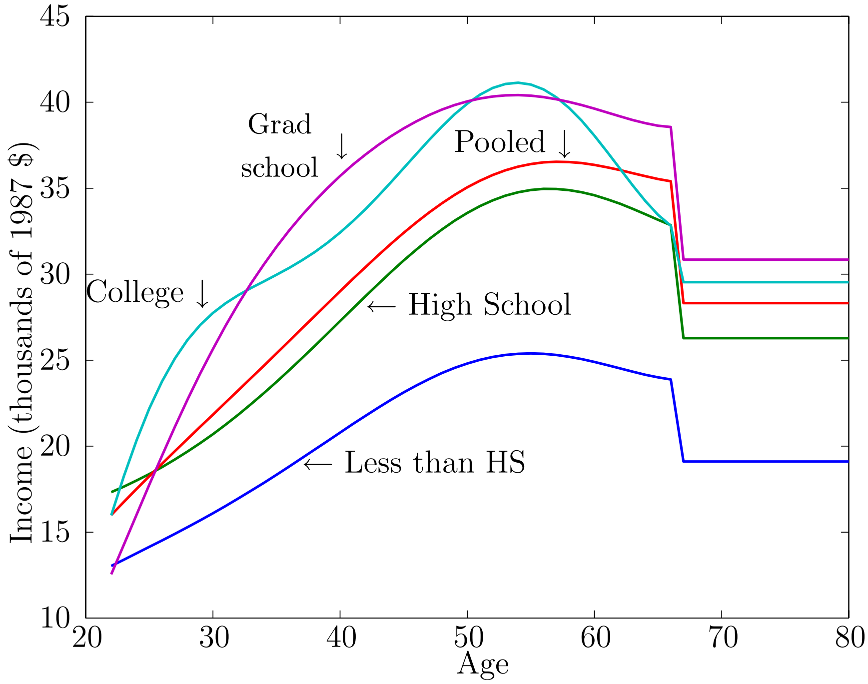

Figure 2: Life-cycle equity holdings at various borrowing rates. Baseline parameter settings for a household from the pooled sample and no labor income risk. `[figa1]`

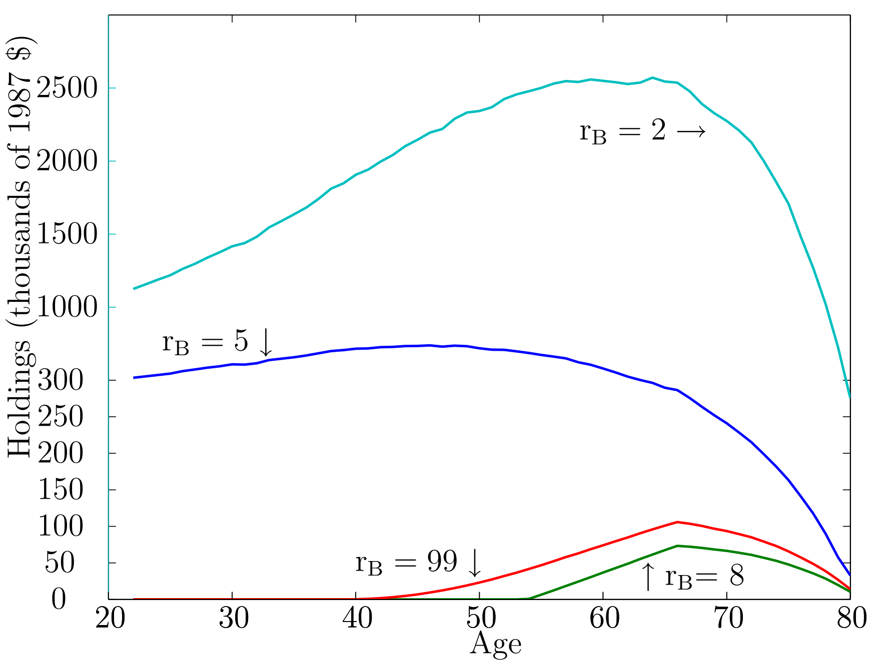

Figure 3: Life-cycle borrowing at various borrowing rates. Baseline parameter settings for a household from the pooled sample and no labor income risk. `[figa2]`

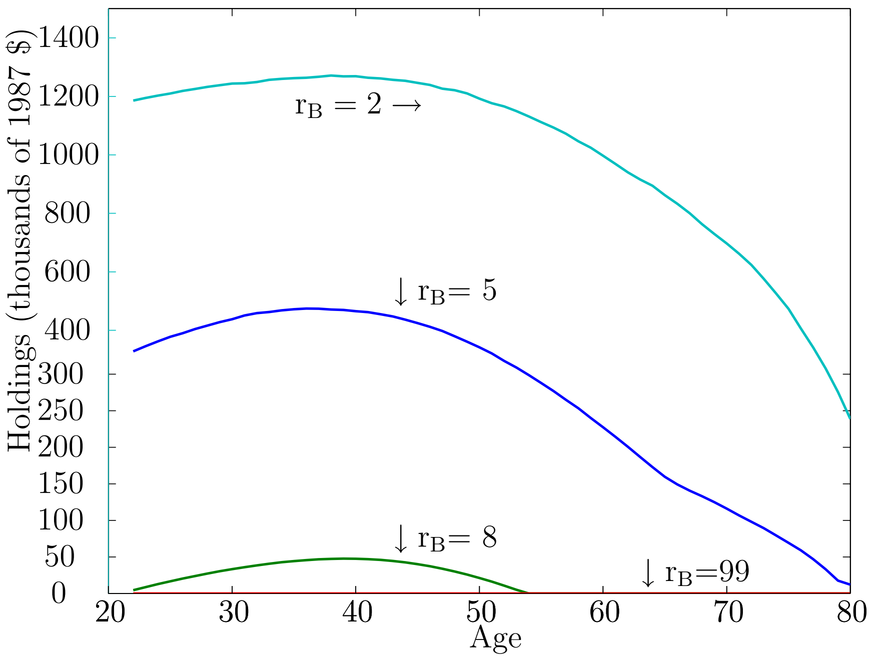

Figure 4: Life-cycle liquid wealth at various borrowing rates. Baseline parameter settings for a household from the pooled sample and no labor income risk. `[figa3]`

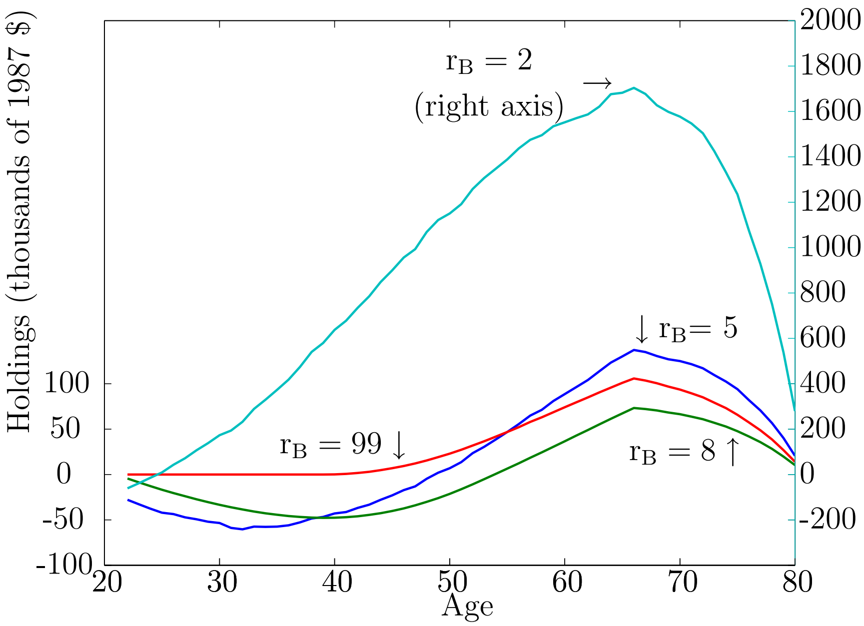

Figure 5: Average equity demand as a function of borrowing rate. Baseline parameter settings for a household from the pooled sample. [figb1]

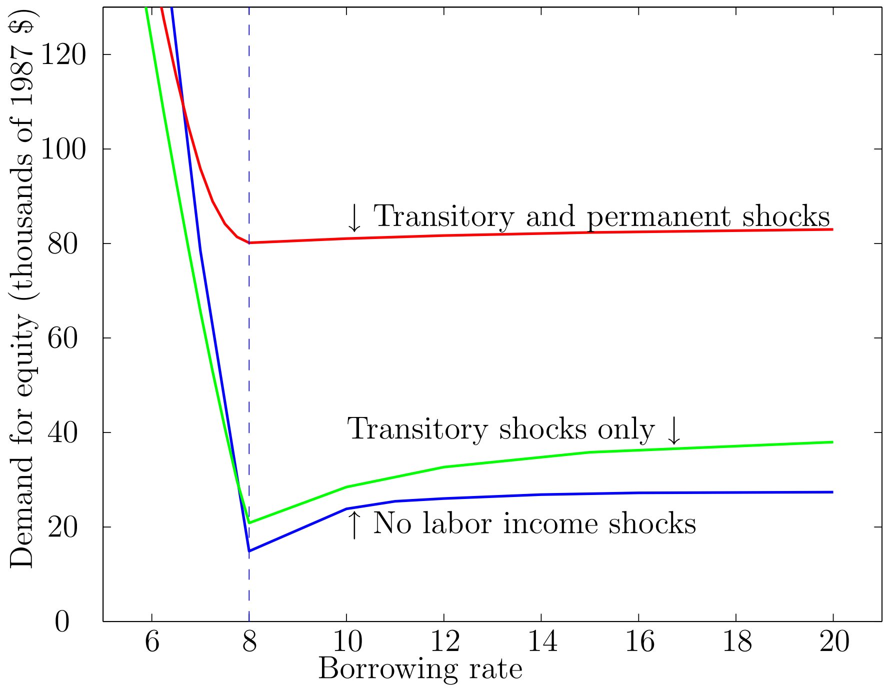

Figure 6: Life-cycle equity holdings with and without labor income risks using pooled sample and baseline parameter settings. [figa6]

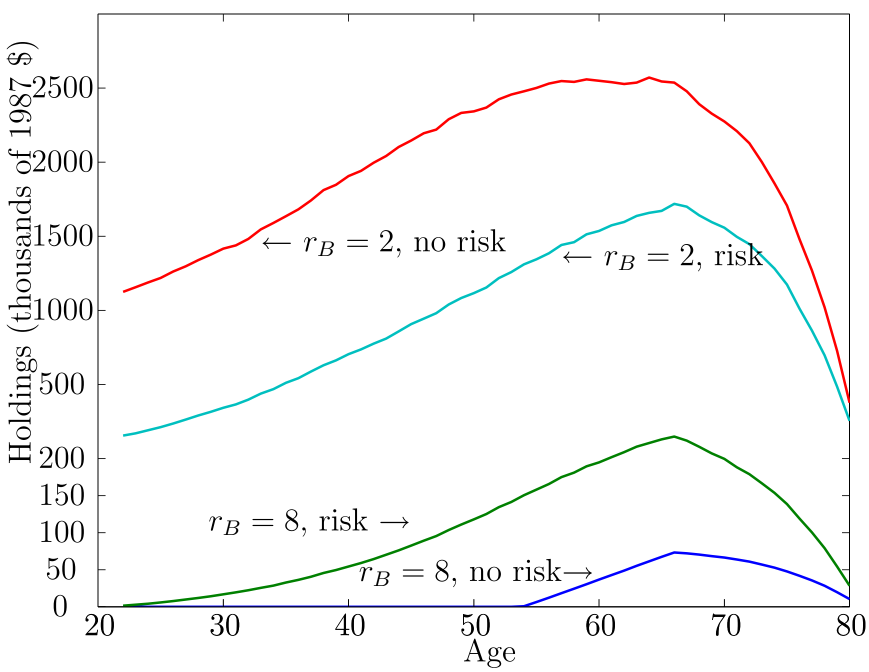

Figure 7: Life-cycle equity holdings at alternative RRA values for pooled sample with risky labor income and baseline parameter settings. [figa4]

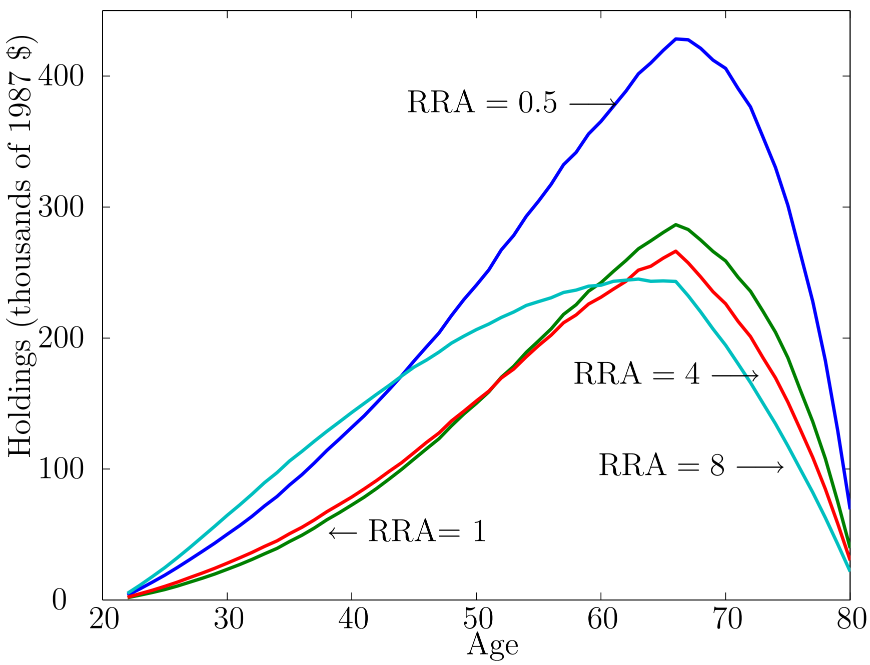

Figure 8: Demand for equity as a function of relative risk aversion. Pooled sample with risky labor income and baseline parameter settings. [figb2]

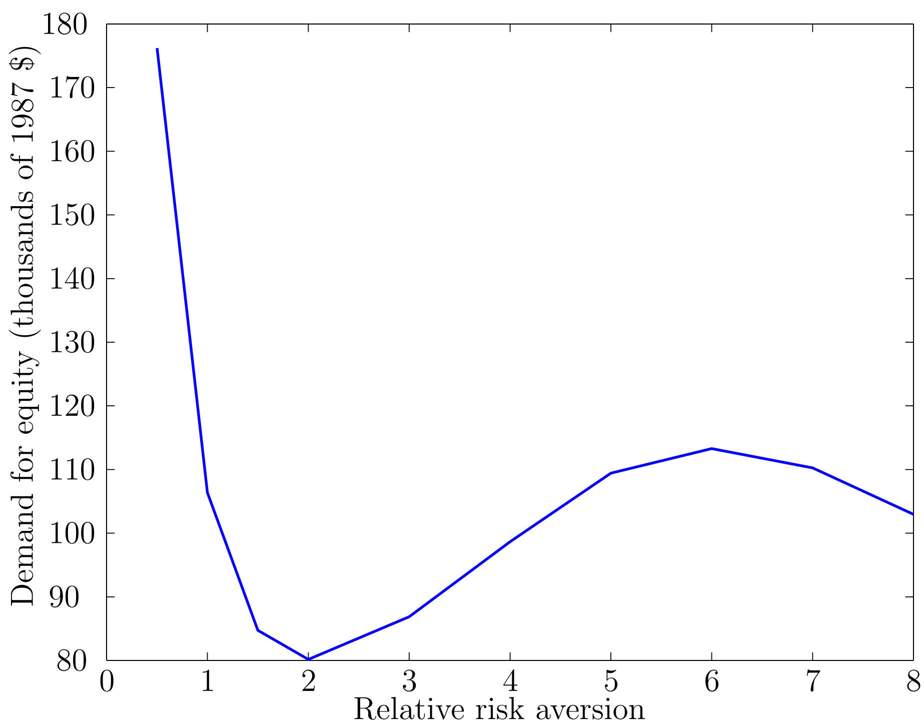

Figure 9: Life-cycle equity holdings for pooled sample and flat income profiles. [figa9]

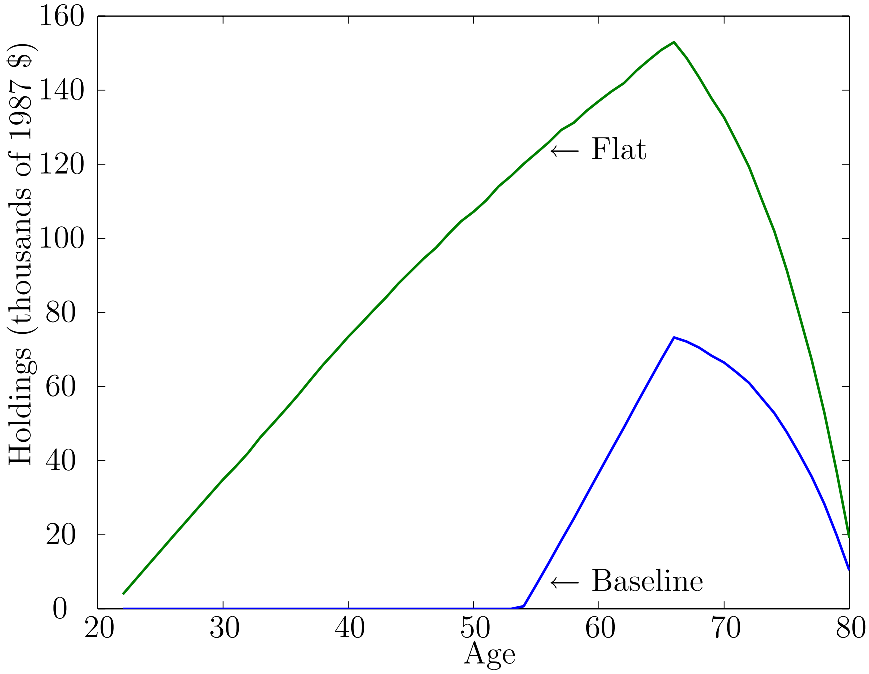

Figure 10: Life-cycle equity holdings for alternative income replacement rates during retirement. [figa8]

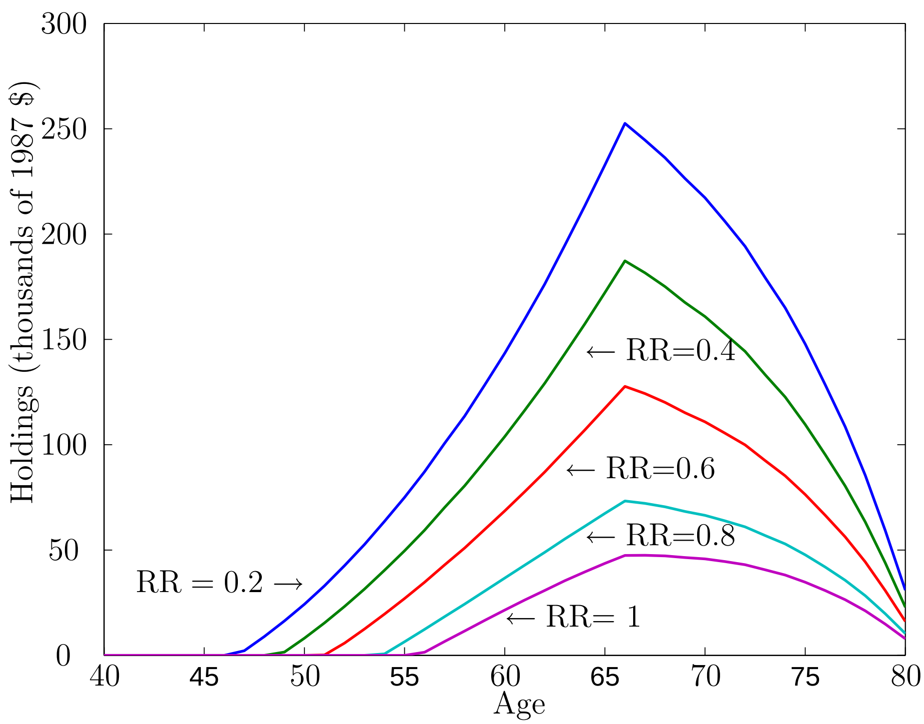

Figure 11: Maximum borrowing over the life cycle as a function of the borrowing rate for baseline specifications with no labor income shocks. [figb4]

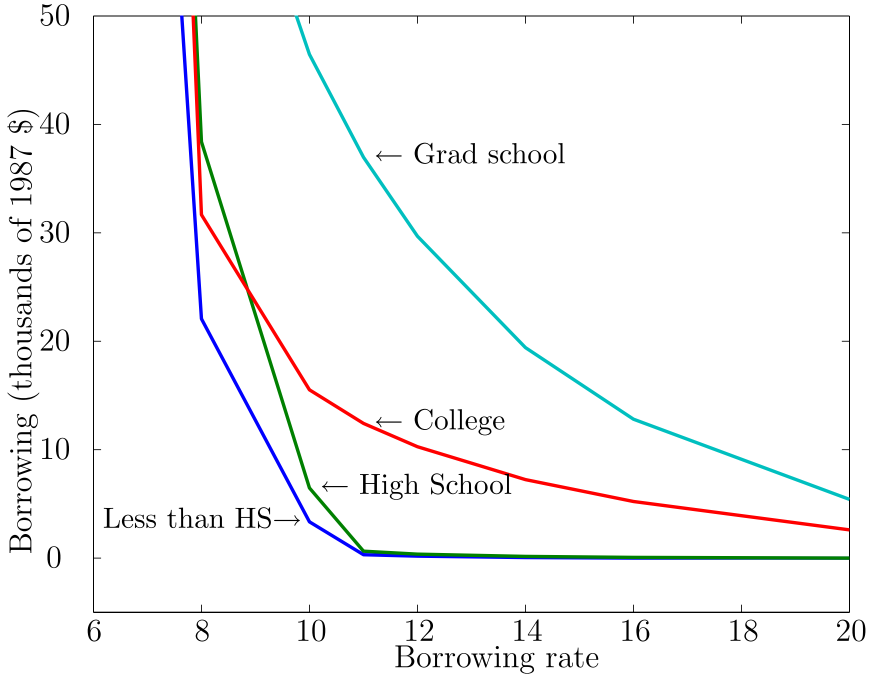

Figure 12: Age that median household enters equity markets for baseline specification with no labor income shocks. [figb5]