Working Papers

Federal Reserve Bank of Boston

No. 05-7

Borrowing Costs and the Demand for Equity over the Life Cycle

Steven J. Davis, Felix Kubler, and Paul Willen

Abstract:
We construct a life-cycle model that delivers realistic behavior for both equity holdings and borrowings. The key model ingredient is a wedge between the cost of borrowing and the risk-free investment return. Borrowing can either raise or lower equity demand, depending on the cost of borrowing. A borrowing rate equal to the expected return on equity — which we show roughly matches the data — minimizes the demand for equity. Alternative models with no borrowing or limited borrowing at the risk-free rate cannot simultaneously fit empirical evidence on borrowing and equity holdings.

JEL Classifications: D91, G11

Steven J. Davis is William H. Abbott Professor of International Business and Economics at the University of Chicago Graduate School of Business and a research fellow at the NBER; Felix Kubler is Professor of Economics at the University of Mannheim; and Paul Willen is a Senior Economist at the Federal Reserve Bank of Boston and a research fellow at the NBER. Their email addresses are steve.davis@gsb.uchicago.edu, fkubler@rumms.uni-mannheim.de, 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/wp/index.htm.

We have benefited greatly from comments by John Heaton, Rick Green, Anna Lusardi, Dave Backus, our discussants, Amir Yaron, Muhammet Guvenen, and Francisco Gomes, two referees, the editor, and many seminar and conference participants. A special thanks to Anil Kashyap for several helpful discussions. We also thank Jonathan Parker for providing parameters of the income processes in Gourinchas and Parker (2002), Annette Vissing-Jorgenson for statistics on the incidence of credit constraints, 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. Davis and Willen gratefully acknowledge research support from the Graduate School of Business at the University of Chicago

This version: May 2, 2005.

1 Introduction

Borrowing presents a problem for life-cycle models of consumption and portfolio choice. In the classic Merton-Samuelson model, modified to include a realistic process for labor income, unsecured borrowing leads to huge, highly leveraged equity positions. For example, with relative risk aversion of 2 and standard specifications for income and asset returns, the model yields average equity holdings more than 20 times bigger than average annual income. To be sure, life-cycle models that preclude borrowing can generate realistic equity holdings, but they fly in the face of evidence that unsecured consumer credit is widely available and widely used. In fact, unsecured debt is much more prevalent than equity in the portfolios of younger households.

In this paper, we construct a life-cycle model that resolves the tension between borrowing and equity holdings. Households can borrow in our setup – but at rates that exceed the risk-free investment return. Given realistic borrowing costs, the model yields both debt positions and equity holdings that fit the main features of the data.

Except for its treatment of borrowing, our preferred model is entirely standard. Agents have time-separable, isoelastic preferences with moderate risk aversion. They face realistic income processes and can invest in risky and risk-free assets. We do not rely on habit formation, self-control problems, myopia or costs of participating and trading in equity markets to obtain sensible life-cycle profiles for borrowing and equity holdings. Neither do we rely on informational barriers, time-varying asset returns or enforcement problems in loan markets. Instead, the key elements of our analysis are realistic borrowing costs and the life-cycle structure. But, as we explain, realistic borrowing costs magnify the impact of certain other frictions – such as fixed costs of participating in equity markets or liquidity benefits from bond holdings – on participation rates and portfolio shares.

Table 1 reports data on the size of the wedge between borrowing costs and the risk-free return. 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 of this wedge 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 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.

The relationship between equity holdings and the cost of borrowing is non-monotonic in our model. To see why, suppose initially that the borrowing rate equals the expected return on equity. No one borrows to buy equity in this case, because the net return is zero and the investment would increase risk exposure. At a slightly lower borrowing rate, however, the net return is positive and the household adopts a small debt-financed equity position. Further reductions in the borrowing rate lead to greater leverage and further increases in equity demand. Now, move in the other direction and consider a borrowing rate that slightly exceeds the equity return. In this case, households with debt hold no equity (because debt repayment offers a better return), so the borrowing rate has no immediate impact on their equity demand. But higher borrowing rates discourage borrowing for consumption-smoothing purposes. As a result, households borrow less at each age, achieve a positive financial position earlier in life, invest in equity at an earlier age and hold more equity at later ages. Further increases in the borrowing rate imply a further upward shift in the life-cycle equity profile, and sufficiently high borrowing costs choke off all borrowing. Hence, equity holdings and participation rates are minimized when the borrowing rate equals the expected return on equity – a scenario consistent with Table 1.

We also develop several other points. First, our model implies high non-participation rates in equity markets, much higher than in otherwise identical models with no borrowing and much closer to the data. Second, even a small wedge between borrowing rates and the risk-free return dramatically reduces the demand for equity. Third, greater income uncertainty raises equity demand in our model with realistic borrowing costs, contrary to its effect in the standard model with no wedge. Fourth, equity demand is a non-monotonic function of relative risk aversion with realistic borrowing costs, again contrary to the standard model. Fifth, and not surprisingly in light of our previous remarks, equity demand is sensitive to the shape of the life-cycle income profile in our preferred model. Finally, we also consider a model with limited borrowing at the risk-free rate and show that it does a poor job of resolving the tension between borrowing and equity holdings. The limited-borrowing model implies that households borrow to finance equity holdings and always exhaust borrowing capacity. Both implications are sharply at odds with observed behavior.

We reiterate that our main goal is to construct a model that delivers realistic life-cycle behavior for both equity holdings and unsecured borrowing. We largely meet that goal, but gaps between theory and data remain. When fit to the evidence on unsecured borrowing and the historical equity premium, equity holdings in our baseline specification are somewhat larger than in the data. And, like other life-cycle models with no liquidity motive for bond holdings, our model does not generate realistic bond portfolio shares with moderately risk-averse investors.

The paper proceeds as follows. The balance of the introduction discusses related research and reviews some important facts about borrowing and equity holdings over the life cycle. Sections 2 and 3 describe the model and choice of parameters. Section 4 considers life-cycle behavior in our preferred model and alternatives, and Section 5 compares model implications with empirical evidence. Where the models fail to fit the facts, we assess the significance of the failures. Section 6 offers some concluding remarks, and an appendix describes our numerical solution method.

1.1 Relationship to the theoretical literature

The structure of our model departs modestly from the seminal work on life-cycle portfolio behavior by Merton (1969) and Samuelson (1969). Indeed, our model differs from Samuelson's discrete-time setup in only three respects: the wedge between borrowing costs and risk-free returns, the presence of undiversifiable income shocks, and the use of realistic income profiles. The wedge and the undiversifiable shocks necessitate a computational approach to the analysis, which we pursue using the same methods as in Judd, Kubler, and Schmedders (2002). In our model, unlike Brennan's

(1971) or Heaton and Lucas's (1997), 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.[1]

Bisin and Gottardi (1991) and Dubey et al. (2003) consider models of adverse selection that endogenously generate differences in prices for buyers and sellers of financial assets. These models can deliver differential borrowing and lending rates, but they do not account for the wedge measured in the last two rows of Table 1, which nets out uncollected loan obligations in order to highlight the cost of producing consumer credit. We take this cost as given and develop its implications for borrowing, equity demand and participation behavior. Why the cost of producing consumer credit is so high is an interesting question that we leave for another occasion.

In order to keep the focus on unsecured borrowing, our model omits ingredients that are probably important for a complete understanding of life-cycle consumption and portfolio behavior. In particular, we omit housing consumption and borrowing secured by housing. Cocco (2004) and Yao and Zhang (2005) argue that a realistic treatment of housing can bring life-cycle models closer to the data. We also ignore the possibility that bonds provide important liquidity services, as argued by Bansal and Coleman (1996) and others.

1.2 Facts about borrowing and equity over the life cycle

Two well-documented sets of facts are relevant to an assessment of our model and alternatives. First, a large percentage of households hold little or no equity. Only 44 percent of households held stock in 1994, a big increase over the 28 percent figure for 1984 (Vissing-Jorgenson, 2002). Participation rates rise with age (Poterba and

[1]Several recent studies analyze consumption and portfolio choice in life-cycle and infinite-horizon models with hard borrowing limits. These studies are cited below or in earlier versions of this paper. We recently became aware of a study by Cocco et al. (2005) that shares several elements of our analysis, including realistic borrowing costs.

Samwick, 2001), education, and income (Mankiw and Zeldes, 1991; Brav and Geczy, 1995), and self-employed persons are more likely to hold stock (Heaton and Lucas, 2000a). To a large degree, low equity market participation can be traced to the fact that many households have little or no financial wealth (Lusardi et al., 2001). Among households that do own equity, most have modest holdings. Vissing-Jorgensen 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.

Second, unsecured consumer credit is widely available and widely used. Durkin (2000, Table 1) reports that 74 percent of all American families had at least one credit card in 1995, and 44 percent of all families had a positive balance after the most recent payment. Despite the high borrowing costs documented in Table 1, many households, especially younger ones, take on substantial unsecured debt. Table 2 provides evidence on this point, showing that many households adopt large debt positions (relative to annual income), and that debt-income ratios decline with age.

Table 2 also reports unused credit as a percent of annual income. The reported measure is a lower bound, because it does not account for the ability to acquire extra credit cards, raise the credit line on existing cards or obtain other forms of personal credit. Most households have unused borrowing capacity, and middle-aged and older households in particular have considerable unused borrowing capacity. This pattern fits with much previous research that finds a declining incidence of binding borrowing constraints with age (for example, Jappelli, 1990 and Duca and Rosenthal, 1993). For a detailed description of life-cycle and cross-sectional variation in household financial positions based on the 1998 SCF, see Kennickell et al. (2000).

2 The model

We consider an optimizing 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 ($\tilde{\varepsilon}_t$):

$$\tilde{y}_t = d_t + \tilde{\eta}_t + \tilde{\varepsilon}_t. \tag{1}$$

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 labor income – consistent with Social Security and most defined benefit pension plans. 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.

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. Households cannot die in debt, which implies that net indebtedness cannot exceed the present value of the household's lowest possible future income stream discounted at r_B. This debt limit is the only constraint on borrowing in our preferred model, but we also consider models that limit borrowing to BL times annual income.

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

$$U(c_t) + \mathrm{E}_t \sum_{a=t+1}^{T} \beta^{a-t} U(\tilde{c}_a) \tag{2}$$

subject to a sequence of budget constraints and possibly a borrowing limit BL, 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.

3 Parameter settings and discretization

Tables 3 summarizes our parameter settings. We set the coefficient of relative risk aversion to 2 in our baseline specification and consider other values ranging from 0.5 to 9. Following 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. We set the correlation of equity returns and labor income shocks to zero.[2] In line with Table 1, we set the baseline borrowing rate to 8 percent, but we also consider a wide range of other values. According to Table 2, more than 10 percent of households with heads under 30 borrow in excess of their annual income, and many other households could borrow similarly large amounts. In this light, we set $BL = 1$ in the model with limited borrowing at the risk-free rate. For the model with no borrowing, $BL = 0$.

For the life-cycle income process, we adopt parameter values estimated by Gourinchas and Parker (2002) from the Consumer Expenditure Survey (CEX) and the Panel Study of Income Dynamics (PSID).[3] 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." (Without these deductions, household income would be about 27 percent higher.) They restrict their sample to male-headed households and attribute the head's age to the entire household.

To estimate the deterministic component of income, GP fit a fifth-order polyno-

[2]Davis and Willen (2000) present evidence of non-zero correlations between labor income shocks and equity returns, and they consider the implications for life-cycle portfolio choice. Haliassos and Michaelides (2003) also study the effect of a non-zero correlation on portfolio choice. Both studies find that correlation values in line with the evidence have modest effects on portfolio choice.

[3]Gourinchas and Parker estimate a life-cycle income process for five education groups and for their full sample, which pools over education groups. To focus on essentials, we restrict attention to their pooled-sample income process. Earlier drafts of this paper report results by education group.

mial in the head's age to CEX data on log family income. To estimate the standard deviation of transitory and permanent income shocks, they 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. The baseline specification in Table 3 reports the standard deviations of the income shocks after adjusting for measurement error.

The resulting expected income profile reflects 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. In the analysis below, we sometimes alter the variances of the income shocks in order to explore how income uncertainty affects equity demand and other outcomes. When we adjust the income process in this way, we also adjust the deterministic income path d_t to preserve the expected income profile.

We select the subjective time discount factor β so that the predicted life-cycle borrowing profile matches the profile in Table 2 as closely as possible. Specifically, given a specification of the income process for our preferred model, we choose β to minimize the average absolute deviation between the mean debt-income ratio in the model and the mean debt-income ratio in Table 2. In computing the average deviation, we weight each age group in proportion to its 1990 U.S. population share. Row 1 of Table 4 shows that a discount factor of 0.933 minimizes the average absolute deviation for our baseline income process. Row 2 carries out the same exercise for the GP income process with no adjustment for measurement error. The greater income uncertainty in Row 2 raises precautionary saving and lowers borrowing, so that a lower discount factor of .914 is needed to match the borrowing profile. Rows 3-5 report the best-fitting discount factors when we turn off one or both income shocks. Overall, the model does a reasonable job of matching the data for each income process. The principal failure relates to borrowing later in the life-cycle. We discuss the fit between the model and the data more extensively in Section 5.

Our model has three sources of randomness: a permanent labor income shock, a transitory income shock, and an asset return shock. We discretize the state space by the method of Tauchen and Hussey (1991), using two points for the permanent shock, two points for the transitory shock, and three points for the asset return shock.[4]

4 The demand for equity over the life cycle

In this section, we explore how equity holdings and other outcomes are affected by four aspects of the household decision problem: (1) the borrowing regime, (2) the shape of the income profile, (3) risk aversion, and (4) undiversifiable income shocks. We also discuss the behavior of bond holdings. Before proceeding, we define some useful terminology.

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.[5] 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.

[4]There is no state of nature with zero income in our discretization. In reality, social safety nets effectively bound income above zero, which argues for a specification with no zero-income state. One might still ask, however, whether our results rely on an overly coarse income grid with a high income floor. To investigate this issue, we experimented with three rather than two grid points for each income shock. It turns out that a finer grid has little impact on model fit; an extra grid point for the permanent shock actually improves the model's fit to the life-cycle profile of equity holdings.

[5]Strictly speaking, the present value of future labor income is a lower bound on true borrowing capacity, which varies with equity holdings. Our numerical solution procedure uses the period-by-period budget constraints, but the concept of borrowing capacity is a useful aid to intuition.

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

4.1 Effect of the borrowing rate and borrowing regime

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 financial 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 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. To illustrate this point, Figure 1 shows life-cycle equity holdings (averaged over many draws) in our baseline specification with alternative borrowing rates. When the borrowing rate equals the risk-free return of 2 percent, households invest enormous amounts in equity throughout the life cycle, a result that is insensitive to the shape of the income profile. Thus, the standard model with $r_B = r_L$ implies equity holdings that dwarf what we see in the data.

A borrowing rate of 5 percent yields much lower equity holdings throughout the life cycle. Why? An increase in the borrowing rate from 2 percent to 5 percent implies a reduction in the leverage premium from 6 percent to 3 percent and a decline in borrowing capacity. The effect on a very young household is easily understood: since it has no financial 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. Indeed, households with a positive leverage premium continue to borrow until the year before death. 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. A household with a 2 percent borrowing rate has much greater wealth at retirement than a household with a 5 percent borrowing rate.

Equity demand behaves differently when the leverage premium is negative. Figure 2 shows that average equity demand *rises* with borrowing costs when the borrowing rate exceeds the return on equity.[6] This result can be understood as follows. When the leverage premium is negative, no household draws on borrowing capacity to buy equity, so that equity demand depends on the level of financial wealth and the share invested in equity. Higher borrowing rates then increase wealth accumulation in two ways. First, they discourage life-cycle consumption smoothing through the loan market, so that households begin accumulating wealth at younger ages. Second, they inhibit reliance on borrowing to smooth transitory income shocks, leading to greater precautionary saving. The first effect involves the shape of the life-cycle expected income profile, and it operates whether or not income is uncertain. The second effect arises from transitory income shocks.

Figure 3 compares the life-cycle pattern of median equity holding in our preferred model with $r_B = 8$ percent to alternative models with no borrowing ($BL = 0$) or limited borrowing at the risk-free rate ($BL = 1$). Both alternatives imply higher equity holdings throughout the life-cycle. The no-borrowing model can be seen as a special case of our preferred model with r_B high enough to choke off all borrowing. Since a borrowing rate equal to the return on equity minimizes the demand for equity, shutting off all borrowing raises equity holdings. The model with limited borrowing at the risk-free rate yields even higher equity holdings, because households adopt a leveraged equity position and exhaust borrowing capacity throughout the life cycle. By exploiting the leverage premium, households accumulate wealth more rapidly, and they invest part or all of this wealth in equity.

[6] In constructing average equity demand from simulated model outcomes, we use population weights for age groups from Bureau of the Census (1994, Table 1).

The model with realistic borrowing costs also implies much higher non-participation rates in equity markets than the alternative models, as seen in Figure 4. In our preferred model with the baseline specification, participation rates are around 25 percent in the first decade of adulthood, and they rise steadily with age to reach 100 percent by age 50. It is worth stressing that this life-cycle participation pattern and the high rates of non-participation do not rest on any friction in the equity market itself. Participation costs, diversification costs, trading costs, and other frictions in the equity market would further reduce participation rates, a point we return to in Section 5.

At a borrowing rate of 8 percent, the median household does not participate in equity markets until age 36 (Figure 3). Higher borrowing costs raise participation rates at all ages. When the interest rate is sufficiently high so as to eliminate borrowing, the median household holds equity at all ages (Figure 3). When faced with a positive leverage premium – as in the model with limited borrowing at the risk-free rate – *every* household holds equity at all ages.

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 noteworthy since the borrowing rates reported in Table 1 lie near estimates of the expected return on equity. Third, households often hold no equity in the model with realistic borrowing costs – in contrast to models with lower borrowing rates in which households always hold equity.

4.2 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 $r_B = r_L$, the shape of the income profile has little effect on equity demand with uncertain labor income and no effect with certain labor income. 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. 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 they attain positive financial wealth. The age at which this occurs depends on the shape of the income profile. The profile shape 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 $r_B \geq \mathrm{E}(\tilde{r}_E)$.

To illustrate the effect of the profile shape, consider the case with $r_B = \mathrm{E}(\tilde{r}_E)$ and no labor income risk. Figure 5 compares life-cycle equity demand in our baseline case with an 80 percent income replacement rate during retirement to three alternatives: a 20 percent replacement rate, a 100 percent replacement rate, and a flat profile with income set to the simple mean of labor income during the working years. The household with a flat profile invests in equity throughout life. 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. A lower replacement rate leads to higher saving, earlier participation in equity markets, and greater equity holdings at each age.

4.3 Effect of undiversifiable labor income risk

How does undiversifiable 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 financial 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.

The first effect dominates when $r_B = r_L$, so that greater income uncertainty lowers equity holdings. The second effect dominates when $r_B = \mathrm{E}(\tilde{r}_E)$. This case differs from the $r_B = r_L$ case for two reasons. First, when $r_B = \mathrm{E}(\tilde{r}_E)$, younger households hold little or no equity. Hence, they cannot offload (much) risk by reducing equity holdings, and the first effect vanishes. Second, it is more costly to rely on borrowing to smooth consumption at a high interest rate, so the precautionary motive for asset accumulation becomes stronger. As a result, income uncertainty increases equity demand when $r_B = \mathrm{E}(\tilde{r}_E)$. In the intermediate case with $r_B \in (r_L, \mathrm{E}(\tilde{r}_E))$, the relation between equity holdings and uncertainty is nonmonotonic, as seen in Figure

6.[7] In unreported results, we verify that the relationships between income uncertainty and equity holdings shown in Figure 6 also hold for lower (RRA=0.5) and higher (RRA=5) values of risk aversion.

To better understand the effects of labor income risk, consider the distinct effects of permanent and transitory shocks on equity market participation rates in our preferred model ($r_B = 8$) and the no-borrowing model ($BL = 0$). Bigger permanent shocks raise precautionary saving and equity holdings in both models. In line with this observation, (unreported) simulations show that a bigger permanent-shock variance leads to higher participation rates in both models. In contrast, transitory income shocks push outcomes away from zero and 100 percent participation. By encouraging precautionary savings, transitory shocks lead to greater equity holdings and higher participation rates. But a sufficiently bad transitory shock (or shock sequence) causes a household to draw down its financial assets and exit the equity market. Thus, 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 asset holdings and turn to borrowing. Figure 7 illustrates these effects of transitory income shocks. Relative to a specification with no income risk, transitory shocks raise participation rates at younger ages and lower them at older ages in both models.

4.4 Effect of risk aversion

Greater risk aversion lowers a household's appetite for risk, and its demand for equity, at a given level of financial wealth. But risk aversion also has a powerful effect on wealth evolution over the life cycle. Higher risk aversion means a higher level of precautionary savings, which raises wealth. Higher risk aversion also means a

[7] Empirical evidence is mixed on the connection between income uncertainty and the demand for risky financial assets. Guiso et al. (1996) find a small, positive relationship between income uncertainty and risky asset shares among Italian households, but Alan (2004) finds little support for a positive relationship in Canadian data. Hochguertel (2003) finds a small positive effect of income uncertainty on risky asset demand in Dutch data, but the effect diminishes or disappears when he allows for unobserved heterogeneity among households. In French data, Arrondel and Masson (2003) find that higher income risk leads to greater holdings of risky financial assets.

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 vary significantly with age and income risk. When the borrowing rate equals the risk-free return, higher risk aversion leads to lower equity holdings throughout the life cycle. With realistic borrowing costs, the story is more complicated.

Figure 8 shows average equity demand as a function of risk aversion in the model with realistic borrowing costs. Absent income risk, higher risk aversion lowers equity demand as in the standard model with equal borrowing and lending rates. But consider specification (2) in Table 4, which uses the unadjusted income variances from Gourinchas and Parker (2002) and a low discount factor. In this example, households are highly impatient and inclined to borrow, but higher risk aversion intensifies the precautionary demand for wealth accumulation. As a result, equity demand rises with risk aversion, until risk aversion is strong enough to yield a portfolio dominated by bonds. Figure 8 also shows that equity demand is a non-monotonic function of the risk aversion parameter for our baseline income process. For relative risk aversion below 2 and above 8, equity demand falls with risk aversion, as predicted by simpler models with $r_B = r_L$ or certain labor income. For relative risk aversion between 2 and 8, equity demand rises with risk aversion. Relative risk aversion near 2 or 3 implies values for equity demand near the (local) minimum.

The effects of risk aversion on participation are similarly ambiguous. Participation rates are high for very low levels of risk aversion ($RRA < 1$) in our baseline specification 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 for non-monotonicity in the level of equity holdings. Gomes and Michaelides (2005) obtain a similar result about the impact of risk aversion on participation, using a life-cycle model with no borrowing, Epstein-Zin preferences, a one-time cost of entry into equity markets, and two risky assets.

4.5 Bond holdings

Given our baseline specification with low risk aversion, households rarely hold bonds in any of the models or borrowing regimes we consider. In this respect, our findings are consistent with previous work in the area by Heaton and Lucas (1997, 2000b), Viciera (2001), Gomes and Michaelides (2005), and Haliassos and Michaelides (2003).

Bodie, Merton, and Samuelson (1992) explain the intuition for low bond shares when labor income shocks are uncorrelated with equity returns. The standard Merton-Samuelson model tells us that a household should invest a fixed fraction of total wealth in risky assets. Total wealth is composed of human wealth and financial wealth. If human wealth is uncorrelated with the risky asset, then it counts toward the bond part of total wealth. The more human wealth a household has, the greater its effective bond position, and the larger the fraction of financial wealth allocated to the risky asset. In our baseline specification, the fraction of human wealth in total wealth almost always exceeds the target fraction of bonds in total wealth. Thus, when possible, households reduce their bond position by borrowing (provided that the borrowing rate is less than the equity return).

Table 5, Panel A shows that households invest exclusively in equities in the baseline parameter specification. When households cannot borrow at the risk-free rate, they invest nothing in bonds, and equity holdings equal financial wealth. When they can borrow at the risk-free rate, they typically do so in order to adopt a leveraged equity position, so that equity holdings exceed net financial wealth.

We can generate positive bond holdings in any of the models by increasing the risk aversion parameter. Lower income replacement rates in retirement also increase the propensity to hold bonds. Panel B in Table 5 provides an illustration by altering these two parameters in the baseline specification. First, we set risk aversion to 6 (compared with 2 in the baseline specification), increasing the desired fraction of total wealth invested in bonds. Second, we lower the replacement rate from 0.8 to 0.2, reducing the value of human wealth. The portfolio share invested in bonds rises with age to offset the life-cycle decline in human wealth.

5 Comparing the models to the data

In this section, we assess four models in relationship to evidence on borrowing, equity holdings, and equity participation rates over the life cycle. The four models are the standard one with unlimited borrowing at the risk-free rate, a model with limited borrowing at the risk-free rate ($BL = 1$), a model with no borrowing ($BL = 0$), and our preferred model with realistic borrowing costs ($r_B = 8$ percent). Our preferred model outperforms the other models in two respects. First, it is the only one that can simultaneously deliver realistic life-cycle profiles for borrowing and equity holdings. Second, the welfare costs of the gap between theoretical predictions and evidence are smallest for our preferred model. We conclude this section with a brief discussion of how margin loans would affect our analysis.

5.1 Realistic borrowing and equity demand

Table 6 shows debt positions, equity holdings, and participation rates over the life-cycle for the four models and in the SCF data. The model with unlimited borrowing at the risk-free rate ($r_B = 2$ percent) produces equity holdings and borrowings that are an order of magnitude greater than in the data. This model cannot be made to fit the data by assuming greater patience or lower income risk, because households will continue to leverage up in the equity market. Nor can reasonable levels of risk aversion fit the data. Even with relative risk aversion of 5, for example, the model predicts borrowing 20 times greater than in the data, and equity holding 10 times greater.

The limited-borrowing model ($r_B = 2$ percent, $BL = 1$) produces outcomes that correspond more closely to the data, but it also fails in several respects. First, it implies much more borrowing than seen in the data. As before, greater patience does not help, because patient households still exploit the equity premium. In fact, willingness to postpone consumption frees up borrowing capacity for investment purposes and leads to even bigger equity holdings. Second, the limited-borrowing model cannot replicate the life-cycle profile of the debt-income ratio unless we vary the exogenous borrowing limit in line with the age profile in the data. The model would still fail to match the evidence in Section 1.2 that unused credit rises with age. Third, the limited-borrowing model predicts 100 percent participation in equity markets at all ages, with equity financed in part by debt. In the data, however, a large fraction of households hold no equity, and few households hold both equity and unsecured debt.

The no-borrowing model and our preferred model with realistic borrowing costs produce similar levels of equity holdings that are much closer to the data. The model with realistic borrowing costs performs better in two key respects. First and foremost, the no-borrowing model is at odds with the prevalence of unsecured borrowing in the data and the widespread availability of unused credit (Table 2). In contrast, our preferred model generates realistic borrowing behavior when calibrated to evidence on the cost of borrowing. Second, our preferred model delivers much higher non-participation rates in equity markets (Figures 4 and 7) and a better fit to equity holdings (Table 6). Our preferred model predicts that a majority of households under the age of forty hold no equity, as in the data, but the no-borrowing model predicts that almost 90 percent of households under forty hold equity. Our preferred model also predicts low equity holdings for younger households, in line with the data, and less than half the levels predicted by the no-borrowing model. In short, a realistic treatment of borrowing also brings the theory closer to the evidence on life-cycle patterns in equity holdings and participation rates.

In terms of explaining the life-cycle behavior of consumption, borrowing, and asset accumulation, housing is the most important ingredient missing from our analysis. A full treatment of housing is beyond the scope of this article, but we can easily compare equity holdings in our models to risky asset holdings in the data, as measured by the sum of equity holdings and housing wealth. In this respect, a comparison of Tables 2 and 7 shows that risky asset holdings (equities) in our preferred model ($r_B = 8$ percent) and in the no-borrowing model are lower than equities plus housing in the data. The gap between theory and evidence for the asset-to-income ratio widens with age. An important topic for future research is the integration of realistic borrowing costs into life-cycle portfolio models that explicitly model housing wealth and consumption.

5.2 Welfare analysis of model failures

Section 4.5 shows that our preferred model, like the alternatives, fails to match evidence on bond holdings. The model also predicts higher equity market participation rates than in the data. How serious are these failures? One useful way to address this question is to quantify the certainty-equivalent consumption cost of deviations between the data and the optimal behavior implied by the model.

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

$$\sum_{t=0}^{T} \beta^t \frac{\overline{c}^{1-\gamma}}{1-\gamma} = U \quad \text{for} \quad \overline{c} = \Big[\frac{1-\gamma}{\sum_{t=0}^{T} \beta^t} U\Big]^{\gamma-1}, \tag{3}$$

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

To measure the costs of suboptimal behavior, we consider three experiments: households do not hold equity, households hold no equity before age 50, and households allocate 50 cents out of every dollar of investment to bonds. We reach two sets of conclusions. First, in our preferred model, the costs of these deviations from optimal behavior are quite small, ranging from .1 to .8 percent of lifetime consumption. Second, the costs are higher for the other models and, in the case of the model with unlimited borrowing at the risk-free rate, dramatically so.

Table 7 shows the results. Panel A considers the baseline specification, and Panel B considers a lower equity return of 6 percent. Two observations motivate a lower equity return. First, many believe that an ex-ante equity return of 8 percent is simply too high. Second, the cost of achieving a diversified equity portfolio lowers the net return, and for most investors mutual funds offer the only feasible means to obtain a broadly diversified portfolio. According to McGrattan and Prescott (2003), mutual fund costs range from 1.3 to 2.5 percent of assets per year in the period from 1980 to 2001.

For the baseline specification, our preferred model implies that the cost of never holding equity is .8 percent of lifetime consumption. The cost of a 50-50 bond-equity mix amounts to .6 percent of consumption. And if the household merely delays equity participation until age 50, the cost amounts to .3 percent of consumption. The costs are lower yet at a 6 percent equity return, as seen in Panel B. For example, at a 6 percent return on equity, waiting till age 50 to participate in equity markets lowers certainty-equivalent consumption by .1 percent, which amounts to $20 per year in 1987 dollars.

The costs are bigger for the other models. The no-borrowing model implies that a no-equity restriction reduces certainty-equivalent consumption by 1.39 percent, nearly three-quarters bigger than in the preferred model. For the limited-borrowing model, the cost of the no-equity restriction is nearly 4 percent of consumption. Finally, the model with unlimited borrowing at the risk-free rate implies enormous welfare costs for suboptimal behavior: a household that refuses to hold equity accepts a 20 percent reduction in lifetime consumption according to this model. It is hard to imagine participation or transactions costs that would overcome a 20 percent or even a 4 percent loss of consumption.

It is useful to assess these results in light of Vissing-Jorgenson's (2002) study of stock market participation costs. She provides an informative discussion of these costs, and she estimates their effects on equity market participation rates and portfolio shares. Based on an after-tax equity premium of 5.6 percent, her estimates imply that a participation cost of $30 per year (in 1984 dollars) is sufficient to account for half of all nonparticipating households, and an annual cost of $175 is sufficient to account for 75 percent. While Vissing-Jorgenson takes low asset holdings as given, our analysis explains low financial wealth as a natural consequence of life-cycle factors and realistic borrowing costs.

Our model abstracts from many real-world features that generate demand for bonds such as participation, diversification, and rebalancing costs, a desire for liquidity, information costs, and so on. Since the gains to holding equity are modest in our preferred model, and very small for a large fraction of households, there is ample scope for these features to reduce equity market participation rates and increase bond portfolio shares.[8] Haliassos and Michaelides (2003) make an identical point in the context of an infinite-horizon model with no borrowing. Our analysis shows that this point carries even greater force in a life-cycle model with realistic borrowing costs than in a model with no borrowing.

5.3 Margin loans

Some commentators have suggested that our results on equity demand and equity market participation would not survive the introduction of margin loans. However, a few observations make clear why the introduction of margin loans would not greatly affect our results. First, initial margin requirements on equity are 50 percent or higher. Thus, for a household with one thousand dollars in financial wealth, a margin loan allows for an equity position of no more than two thousand dollars. Second, the data show a large wedge between margin loan rates and risk-free returns. Kubler and Willen (2002) report that as of July 8, 2002, the rates on margin loans of less than $50,000 at five major brokerage houses (The Vanguard Group, Fidelity Investments, Charles Schwab, Salomon Smith Barney, and UBS Paine Webber) exceed the rate on 90-day U.S. Treasury Bills by 357 to 570 basis points, depending on brokerage house and loan size. Even at these rates, brokerage houses require credit checks and reserve the right to deny margin credit or impose higher margin rates. Finally, the combination of unsecured borrowing and margin loans does not offer an attractive leverage premium. For example, at an 8 percent expected return on equity, a risk-free rate of 1.68 percent (the return on 90-day U.S Treasury Bills as of July 8, 2002) and a 4.63 percent margin loan premium, the expected return on a margin-leveraged equity portfolio is $(1/.5)8 - (1.68 + 4.63) = 9.69$ percent. Combined with a wedge of 7.5 percentage points on unsecured borrowing, roughly the midpoint of the Table 1 values,

[8] Certain frictions (for example, a fixed cost of equity holding) lower equity market participation but do not raise bond shares conditional on participation. Other frictions (for example, proportional trading costs in equity markets) also raise bond shares conditional on participation. See Aiyagari and Gertler (1991) for an early analysis of trading frictions in equity markets.

the fully leveraged portfolio offers a leverage premium of $9.69 - (1.68 + 7.5) = .51$ percent. That is, the fully leveraged portfolio offers an expected return premium of 51 basis points with a standard deviation of $2 \times 15 = 30$ percent. At a 6 percent expected return on equity, the margin-leveraged equity portfolio yields a negative return.

6 Concluding Remarks

We showed that a model with a wedge between borrowing costs and the risk-free investment return can simultaneously deliver sensible life-cycle profiles for debt and equity holdings and high rates of non-participation in equity markets. Realistic borrowing costs dramatically reduce equity holdings, and equity demand is at its minimum when the borrowing rate equals the expected return on equity. The model with realistic borrowing costs does a better job of fitting observed life-cycle patterns in borrowing, equity market participation, and equity accumulation than alternative models with no borrowing or limited borrowing at the risk-free rate.

The opportunity to borrow at realistic rates in a life-cycle setting has important consequences for wealth accumulation. Because households face an upward-sloping income profile, they borrow in the early part of the life cycle, which delays the age at which they participate in equity markets or accumulate significant holdings. This implication of our model helps explain the low equity holdings of most households in the face of an apparently high equity premium. Heaton and Lucas (2000b), Attanasio, Banks, and Tanner (2002) and others have emphasized this aspect of household behavior as an important puzzle.

Our analysis points to several directions for future research. We mention two here. First, our model implies that most households accumulate little or no financial wealth until middle age, consistent with much empirical evidence (for example, Kennickell et al., 2000, and Lusardi et al., 2001). Given its simplicity and its assumption of time-consistent, rational consumers, our model and analysis challenge claims that households save too little, or that they should be prompted to save more. A natural next step is to enrich the model to account for housing consumption and real estate wealth and for the liquidity benefits of safe assets. We plan to evaluate richer versions of the model against a number of facts about consumption, home ownership, wealth accumulation, and portfolio behavior over the life cycle.

Second, our analysis highlights the role of borrowing costs and leverage as key factors in the demand for risky assets. While margin loans provide limited scope for leveraged equity holdings, as we have shown, corporate bonds, government securities, real estate, and small business wealth are often subject to much less stringent restrictions on leverage. Kubler and Willen (2002) consider a richer version of our model to address portfolio choice in a broader setting that encompasses a fuller menu of risky assets and leveraging methods. Leverage characteristics turn out to have important implications for portfolio shares, but the cost and availability of unsecured borrowing continue to play a central role in the demand for risky assets.

Appendix: Computational details

Our problem can be formulated as a standard finite-horizon dynamic program and solved by backward induction. It is well known (see Judd (1998), Chapter 12) that this type of problem can be solved by value-function or policy-function iteration. For reasons that will become clear, we use policy-function iteration. At age 80, the solution is trivial: consume everything. We then solve for optimal consumption and portfolio choice at age 79, conditional on financial wealth, income, the state of the world, and the (degenerate) policy rule at age 80. Next, we solve for consumption and portfolio choice at age 78 – again conditional on financial wealth, income, the state of the world, and the calculated optimal policy rule at age 79. And so on.

The problem therefore reduces to solving two-period optimization problems and to approximating policy rules as functions of a minimal set of state variables (including current age). To model the wedge between borrowing and lending rate, we introduce two different bonds: a "lending-bond" that cannot be sold short and a "borrowing-bond" that can only be sold short. Thus, we have three assets in the model and short-sale constraints on all three. The short-sale constraints frequently bind, and they create non-differentiabilities in the policy function at points in the state space where they just become binding. The two-period optimization problem, even though relatively small, is therefore not completely trivial. For smooth, convex problems it is usually easy to solve the optimization problem by steepest descent methods (Judd (1998), Chapter 4). However, constraints on choice sets require the use of constrained optimization routines that are often numerically unstable and costly to set up. Since we need to solve the two-period problem at each iteration (age) at many points in the state space, it is essential to employ reliable methods with low set-up costs. We found that standard constrained-optimization routines (for example, NPSOL or MINOS) involved high setup costs or lacked sufficient reliability.

An alternative is to discretize both the state space and the space of available policies and use simple grid search methods to find the optimal policy at each iteration. The advantage of this method is numerical stability; the disadvantage is that it is very slow. In particular, since we effectively have three assets, a grid search approach is not feasible if we want to obtain sufficiently accurate solutions. Instead, we exploit the fact that the first-order conditions are necessary and sufficient and solve them, at each iteration, as a nonlinear system of equations. This is also why we use policy-function iteration. Since we work with the first-order conditions, only the derivative of the value function enters the two-period problem. Since we do not need the value function in our analysis, it is easier to work with the policy function directly.

We now describe in detail two aspects of the solution procedure: how to reduce the endogenous state space to one variable, and how to solve the two-period problem effectively. First, we use preference homotheticity to simplify the problem and reduce the number of continuous state variables to one. Second, we explain how to solve the two-period optimization problem.

A little notation will help. Let z_t be a Markov chain with finite support $z \in \{1, ..., S\}$ and transition π. Gross equity returns are $\tilde{R}_E(z_t)$, 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 $y(z^t)$ denote income at time t.

Preference homotheticity allows us to simplify the problem by combining wealth and income into one variable (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_t. The optimal policy rule for the next period specifies investment of $F^i_{t+1}(\Xi_{t+1}(z'), y_{t+1}(z'); z')$ 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 financial wealth Ξ.

$$
\begin{aligned}
&\max_{F^L,\ F^B,\ F^E} \quad \frac{c_t^{1-\gamma}}{1-\gamma} + \beta\,\mathrm{E}_t\Big(\frac{c_{t+1}^{1-\gamma}}{1-\gamma}\Big) \qquad (4)\\
&\text{s.t.} \quad c_t = y + \Xi - F^L + F^B - F^E;\\
&c_{t+1}(z') = y(z') + \Xi(z') - F^L(\Xi(z'), y(z'); z') +\\
&\qquad F^B(\Xi(z'), y(z'); z') - F^E(\Xi(z'), y(z'); z'), \quad \forall z' \in \{1, ..., S\};\\
&\Xi(z') = F^L R_L - F^B R_B + F^E R_E(z'), \quad \forall z' \in \{1, ..., S\};\\
&F^L \geq 0, \quad F^B \geq 0, \quad F^E \geq 0;
\end{aligned}
$$

where we suppress time subscripts on variables other than consumption to reduce notational clutter.

Now divide through by y_t, define $x_t = c_t/y_t$, and consider the two-period optimization problem:

$$
\begin{aligned}
&\max_{f^L,\ f^B,\ f^E} \quad \frac{x_t^{1-\gamma}}{1-\gamma} + \beta\,\mathrm{E}_t\Big(\frac{x_{t+1}^{1-\gamma}}{1-\gamma}\Big) \qquad (5)\\
&\text{s.t.} \quad x_t = \xi - f^L + f^B - f^E;\\
&x_{t+1} = \frac{y_{t+1}(z')}{y_t}\big[\xi(z') - f^L(\xi(z'); z') + f^B(\xi(z'); z') - f^E(\xi(z'); z')\big];\\
&\xi(z') = \frac{y_{t+1}(z') + f^L R_L - f^B R_B + f^E R_E(z')}{y_{t+1}(z')};\\
&f^L \geq 0, \quad f^B \geq 0, \quad f^E \geq 0.
\end{aligned}
$$

Observe that the policy rules are now functions of a single endogenous state variable, ξ, the ratio of financial wealth plus current income to current income. This reduction in the dimensionality of the state space greatly simplifies computation. We can recover the solution to the original problem (4) by multiplying the solution to the transformed problem (5) by current income:

$$\begin{gathered} c_t = y_t x_t \\ F^L = y_t f^L, \quad F^B = y_t f^B, \quad F^E = y_t f^E \end{gathered} \tag{6}$$

To solve the transformed two-period problem, we solve the associated Kuhn-Tucker conditions – a nonlinear system of equations and inequalities that is necessary and sufficient for optimality. Following Garcia and Zangwill (1981, pages 65-68), we use a change of variables to eliminate inequalities in the Kuhn-Tucker conditions and state the optimality conditions as a system consisting solely of equations. The resulting system has 3 unknowns corresponding to the three asset holdings. In particular, let $\eta_j \in \Re$ for $j = 1, 2, 3$, and define the Kuhn-Tucker multiplier for asset j, $\mu_j = (\max\{0, -\eta_j\})^3$. The consumer's holding of asset j is $\theta_j = (\max\{0, \eta_j\})^3$. Note that θ and μ are twice continuously differentiable, and that the complementary slackness conditions hold:

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

We implement our solution algorithm using Fortran 90. A simple Newton method usually works well as a nonlinear 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.

Lastly, we draw attention to two practical aspects of our computational solution. First, the range of $f_t^j(\xi; z)$ will generally depend on t and z. In practice, we set arbitrary bounds on the range that vary only with t. We then verify that these bounds never bind in the simulations. Second, in generating $f_t^j(\xi; z)$, we don't solve (5) for every possible value of ξ. Instead, we solve (5) for a finite number of values of ξ and 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 on a Pentium III computer with a 1.2 Ghz processor and 1 GB of RAM 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 return on equity.

References

[1] Aiyagari, S. R., Gertler, M., 1991. Asset returns with transactions costs and uninsured individual risk. *Journal of Monetary Economics* 27, 311-331.

[2] Alan, S., 2004, Precautionary wealth and portfolio allocation: Evidence from Canadian microdata, SEDAP Research Paper no. 117, McMaster University, Canada.

[3] Altig, D., Davis, S., 1992. The timing of intergenerational transfers, tax policy and aggregate saving. *American Economic Review* 82, 1199-1220.

[4] Ameriks, J., Zeldes, S. P., 2001. How do household portfolio shares vary with age? Unpublished working paper. Columbia University.

[5] Arrondel, L. and A. Masson, 2003, Stockholdings in France. In *Stockholdings in Europe,* L. Guiso, M. Haliassos and T. Jappelli, editors. Houndmills, Basingstoke, Hampshire, UK: Palgrave-Macmillan.

[6] Attanasio, O. P., Banks, J., Tanner, S., 2002. Asset holding and consumption volatility. *Journal of Political Economy* 110, 771-792.

[7] Bansal, R. and W. J. Coleman. 1996. A Monetary Explanation of the Equity Premium, Term Premium and Risk-Free Rate Puzzles. *Journal of Political Economy* 104, 1135-1171.

[8] Bisin, A., Gottardi, P., 1991. Competitive equilibria with asymmetric information. *Journal of Economic Theory* 87, 1-48.

[9] Bodie, Z., Merton, R., Samuelson, W., 1992. Labor supply flexibility and portfolio choice in a life-cycle model. *Journal of Economic Dynamics and Control* 16, 427-449.

[10] Brav, A., Geczy, C., 1995. An empirical resurrection of the simple consumption CAPM with power utility. Unpublished working paper, University of Chicago.

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

[12] Bureau of the Census, U.S. Department of Commerce, 1994. *1990 Census of Population: Education in the United States.* U.S. Government Printing Office, Washington, D.C.

[13] Campbell, J. Y., 1999. Asset prices, consumption and the business cycle. In Taylor, J. and Woodford, M. (Eds.), *Handbook of Macroeconomics*, Volume 1C. North-Holland, Amsterdam.

[14] Cocco, J., 2004. Portfolio choice in the presence of housing. *Review of Financial Studies*, forthcoming.

[15] Cocco, J., F. Gomes, and P. J. Maenhout, 2005. Consumption and Portfolio Choice over the Life-Cycle. *Review of Financial Studies*, forthcoming.

[16] Davis S. J. and P. S. Willen, 2000. Using financial assets to hedge labor income risks. Unpublished working paper, University of Chicago.

[17] Deaton, A., 1991. Saving and liquidity constraints. *Econometrica* 59, 1221-1248.

[18] Dubey, P., Geanakoplos, J., Shubik, M., 2003. Default and punishment in general equilibrium. Unpublished working paper, Yale University.

[19] Duca, J. V., Rosenthal, J., 1993. Borrowing constraints, household debt, and racial discrimintion in loan markets. *Journal of Financial Intermediation* 3, 77-103.

[20] Durkin, T. A., 2000. Credit Cards: Use and Consumer Attitudes. *Federal Reserve Bulletin* 86, 623-634.

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

[22] Gomes, F., Michaelides, A., 2005. Optimal life-cycle asset allocation: Understanding the empirical evidence. *Journal of Finance*, 60, no. 2, 869-904.

[23] Gourinchas P.-O., Parker J., 2002. Consumption over the life cycle. *Econometrica* 70, 47-89.

[24] Guiso, L., T. Jappelli and D. Terlizzese, 1996. Income risk, borrowing constraints, and portfolio choice. *American Economic Review*, 86, no. 1 (March), 158-172.

[25] Haliassos, M., Michaelides, A., 2003. Portfolio Choice and Liquidity Constraints. *International Economic Review* 44, 143-177.

[26] Heaton, J., Lucas, D., 1997. Market frictions, savings behavior, and portfolio choice. *Macroeconomic Dynamics* 1, 76-101.

[27] Heaton, J., Lucas, D., 2000a. Portfolio Choice and asset prices: The importance of entrepreneurial risk. *Journal of Finance* 55, 1163-98.

[28] Heaton, J., Lucas, D., 2000b, Portfolio Choice in the Presence of Background Risk. *Economic Journal* 110, 1-26.

[29] Hochguertel, S., 2003, Precautionary motives and portfolio decisions, *Journal of Applied Econometrics*, 18, no. 1 (January/February) , 61-77.

[30] Jappelli, T., 1990. Who Is Credit Constrained in the U.S. Economy? *Quarterly Journal of Economics* 105, 219-234.

[31] Judd, K., 1998, *Numerical Methods in Economics*, MIT Press, Cambridge, MA.

[32] Judd, K., Kubler, F., Schmedders, K., 2002. Computational methods for dynamic equilibria with heterogeneous agents. In Dewatripont, M., Hansen, L. P., Turnovsky, S. J. (Eds.), *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress.* Cambridge University Press, Cambridge, U.K.

[33] Kennickell, A. B., Starr-McCluer, M., Surette, B. J., 2000. Recent changes in U.S. family finances: Results from the 1998 Survey of Consumer Finances. *Federal Reserve Bulletin* 86, 1-29.

[34] Kubler, F. and P. S. Willen, 2002. Margin borrowing and life-cycle portfolio choice. Work in progress.

[35] Lusardi, A., Cossa, R., Krupka, E. L., 2001. Savings of young parents. *Journal of Human Resources* 36, 762-794.

[36] Mankiw, N. G., Zeldes, S. P., 1991. The consumption of stockholders and non-stockholders. *Journal of Financial Economics* 29, 97-112.

[37] McGrattan, E. R., Prescott, E. C., 2003. Average debt and equity returns: Puzzling? *Federal Reserve Bank of Minneapolis Research Department Staff Report* 313.

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

[39] Poterba, J. M., Samwick, A. A., 2001. Portfolio allocations over the life cycle, in Ogura, S. Tachibanaki, T., Wise, D. (Eds.), *Aging Issues in the United States and Japan*. University of Chicago Press, Chicago.

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

[41] Tauchen, G., Hussey, R., 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica* 59, 371-96.

[42] Viciera, L. M., 2001. Optimal portfolio choice for long-horizon investors with nontradable labor income. *Journal of Finance* 56, 433-470.

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

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

[45] Yao, R. and H. Zhang. 2005. Optimal consumption and portfolio choice with risky housing and borrowing constraints. *Review of Financial Studies*, 18, no. 1 (Spring), 197-239.

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

	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: Unsecured Debt, Unused Credit, Stocks and Housing Wealth by Age of Household Head, As Percent of Annual Income

Age Group	Unsecured debt			Unused credit			Equity holdings (Stocks)			Stocks plus home equity		
	Median	90th percent	Mean	Median	90th percent	Mean	Median	90th percent	Mean	Median	90th percent	Mean
23 - 29	19	102	28	12	31	13	0	38	6	2	123	24
30 - 39	12	66	18	15	32	13	0	77	18	37	228	68
40 - 49	9	53	15	18	41	17	3	150	45	85	402	144
50 - 59	5	52	13	21	51	23	7	260	91	154	635	261
60 - 69	0	42	9	25	112	37	0	414	178	250	1399	618
70 - 79	0	25	4	27	135	48	0	716	299	487	2454	1209

Notes:

1. Income is pre-tax household income from wages and salaries plus pre-tax retirement income. Retirement income includes annuities and benefits from Social Security, defined-benefit pensions, and disability programs. It does not include non-annuity income from assets in defined-contribution retirement plans.

2. Unsecured debt is the sum of credit card balances, installment loans, and other debt not secured by real estate, vehicles, etc. Credit card balances are measured after the most recent payment.

3. Unused credit equals the difference between the household's credit limit on its existing credit cards and its actual credit card balance.

4. Equity is the sum of directly held stock and stock held through mutual funds, including investments held in defined-contribution retirement plans. Housing wealth is measured net of mortgage loans.

Source: Authors' calculations from the 1995 and 1998 Survey of Consumer Finances. After deleting households that report neither labor nor retirement income, the sample contains 7,830 observations. We compute all statistics using SCF sampling weights to correct for the over-sampling of households with high net worth. We compute means using trimmed samples that exclude the top 5 percent of the ratio values.

Table 3: Parameter Settings

Parameter	Baseline	Alternative values
Relative risk aversion	2	0.5 to 9
Annual discount factor	.933	.914 to .982
Age of labor force entry	21	
Age of retirement	65	
Age of death	80	
std $(\Delta\tilde{\eta})$ (permanent shock)	12 percent	0, 15 percent
cov $(\Delta\tilde{\eta}, \tilde{r}_E)$	0	
std $(\tilde{\varepsilon})$ (transitory shock)	15 percent	0, 21 percent
cov $(\tilde{\varepsilon}, \tilde{r}_E)$	0	
Replacement rate	80 percent	20, 100 percent
n (for retirement basis)	30	
r_L (risk-free return)	2 percent	
r_B (borrowing rate)	8 percent	2, 5, 6-20 percent
$\mathrm{E}(\tilde{r}_E)$ (equity return)	8 percent	6 percent
std $(\tilde{r}_E)$	15 percent	
Borrowing limit	None	0, 1 times annual income

Table 4: Calibration of subjective time discount factor to the debt-income profile

#	β	std($\Delta\tilde{\eta}$)	std($\tilde{\varepsilon}$)		Means, Percent of Annual Income							Average absolute deviation from data
					23-29	30-39	40-49	50-59	60-69	70-79	23-79	
Data				Debt	28	18	15	13	9	4	15	
				Equity	6	18	45	91	178	299	68	
1	0.933	12	15	Debt	28	20	3	0	0	0	9	7.0
				Equity	3	14	68	187	330	234	128	53.5
2	0.914	15	21	Debt	30	25	7	1	0	2	11	6.8
				Equity	7	18	57	140	201	102	84	38.6
3	0.967	0	0	Debt	33	19	0	0	0	0	9	7.9
				Equity	0	2	83	254	479	404	184	99.0
4	0.912	15	0	Debt	25	19	0	0	0	0	8	7.7
				Equity	0	1	36	115	165	78	62	39.0
5	0.972	0	21	Debt	31	18	2	0	0	0	9	7.1
				Equity	10	41	148	344	607	505	253	160.1

Notes:

1. Debt-income and equity-income ratios in the data are computed from the Survey of Consumer Finances as described in Table 2. The average ratios for ages 23-79 are computed using SCF sample weights.

2. For model specifications 1 through 5, the reported time discount factor β minimizes the mean absolute deviation between the debt-income ratio in the model and the debt-income ratio in the data. The average is taken over the indicated age groups with weights proportional to the 1990 U.S. age distribution, as reported in Table 1 of Bureau of the Census (1994).

3. std($\Delta\tilde{\eta}$) and std($\tilde{\epsilon}$) denote standard deviations of the permanent and transitory income shocks, respectively. The remaining parameters are set to the baseline values reported in Table 3.

3. The rightmost column reports the mean absolute deviation between the model and data for the debt-income and equity-income ratios.

Table 5: Equity and bond holdings over the life cycle in the data and four models

Equity holdings as a percent of gross and net financial wealth

Age Group	Data		Model							
			$r_B = 8\%$		No borrowing		$r_B = 2\%$			
							$BL = 1$		No BL	
	GFW	NFW	GFW	NFW	GFW	NFW	GFW	NFW	GFW	NFW
A. Baseline – RRA=2, RR=80%										
23 - 39	35	75	100	-58	100	100	100	-1262	100	636
40 - 59	45	50	100	101	100	100	100	166	100	232
60 - 79	31	31	100	100	100	100	100	128	100	168
23 - 79	28	29	100	107	100	100	100	166	100	227
B. Alternative – RRA=6, RR=20%										
23 - 39	35	75	100	100	100	100	100	164	100	241
40 - 59	45	50	87	87	87	87	82	88	81	89
60 - 79	31	31	50	50	50	50	45	45	48	48
23 - 79	28	29	59	59	59	59	60	63	63	69

Notes:

1. Gross financial wealth (GFW) is equity plus bonds. Net financial wealth (NFW) is equity plus bonds minus debt. Entries report total equity of households in the group divided by total wealth of the same households. A negative entry under NFW means that net financial wealth is negative, and an entry that exceeds 100 percent means that the average household has a leveraged equity position.
2. Bond holdings as a percent of GFW equal 100 minus the reported figure for equity holdings.
3. Parameters are set to the baseline values reported in Table 3 unless otherwise noted. Averages over 23-79 computed using 1990 Census populations weights as described in Table 4.

Table 6: Debt, equity, and equity market participation over the life cycle

Age Group	Data			Model											
				$r_B = 8$			No borrowing			$r_B = 2$					
										Borrowing Limit= 1			No Borr. Limit		
	Debt	Equity	Ptcp	Debt	Equity	Ptcp	Debt	Equity	Ptcp	Debt	Equity	Ptcp	Debt	Equity	Ptcp
23 - 29	28	6	41	28	3	25	0	15	80	100	77	100	1212	1254	100
30 - 39	18	18	51	20	14	51	0	36	95	100	110	100	1394	1694	100
40 - 49	15	45	56	3	68	93	0	93	100	100	182	100	1561	2204	100
50 - 59	13	91	57	0	187	100	0	205	100	100	299	100	1723	2820	100
60 - 69	9	178	41	0	330	100	0	345	100	100	451	100	1888	3679	100
70 - 79	4	299	34	0	234	100	0	243	100	100	352	100	1196	2867	100
23 - 79	15	68	49	9	128	77	0	146	96	100	234	100	1506	2364	100
Average Absolute deviation from data				7	54	34	15	67	48	85	141	52	1491	2271	52

Notes:

1. Debt-income and equity-income ratios are expressed as percentages of annual income. "Ptcp" is the percent of households with positive equity holdings.

2. In the "Data", we compute averages using SCF sample weights. For the models, we compute averages using 1990 population weights as described in Table 4. We also computed the average absolute deviations between the models and the data using 1990 population weights.

3. All parameters are set to baseline values in Table 3 unless otherwise noted.

Table 7: Welfare costs of restrictions on equity holdings, alternative borrowing regimes

	$r_B = 8$	No borrowing	$r_B = 2$, $BL = 1$	$r_B = 2$, No BL
A. Baseline				
CE consumption in ’000s of $	19.13	19.08	20.08	25.26
No equity				
CE consumption in ’000s of $	18.97	18.82	19.32	20.08
Δ in $	-154	-265	-755	-5181
Δ in %	-0.80	-1.39	-3.76	-20.51
No equity until age 50				
CE consumption in ’000s of $	19.07	18.92	19.50	20.68
Δ in $	-60	-158	-574	-4577
Δ in %	-0.31	-0.83	-2.86	-18.12
50/50 bond equity mix				
CE consumption in ’000s of $	19.01	18.91	19.64	*
Δ in $	-114	-172	-436	*
Δ in %	-0.60	-0.90	-2.17	*
B. $E(\tilde{r}_E) = 6\%$, $\beta = 0.937$				
CE consumption in ’000s of $	19.08	19.01	19.75	22.19
No equity until age 50				
CE consumption in ’000s of $	19.06	18.93	19.44	18.60
Δ in $	-20	-80	-315	-3589
Δ in %	-0.10	-0.42	-1.60	-16.18
50/50 bond equity mix				
CE consumption in ’000s of $	19.02	18.93	19.52	*
Δ in $	-57	-85	-228	*
Δ in %	-0.30	-0.45	-1.15	*

Notes:

1. See text for the calculation of certainty-equivalent (CE) consumption. Consumption is measured in 1987 dollars.

2. Baseline parameter settings except as noted.

3. In the model with $r_B = 2$ and no borrowing limit, the household can circumvent a minimum bond requirement by taking on more debt. If, instead, we require a long position in bonds equal to equity holdings, then the welfare cost of imposing a 50/50 bond-equity mix is very large in the model with $r_B = 2$ and no borrowing limit.

Figure 1: Mean life-cycle equity holdings at various borrowing rates. Baseline parameter settings.

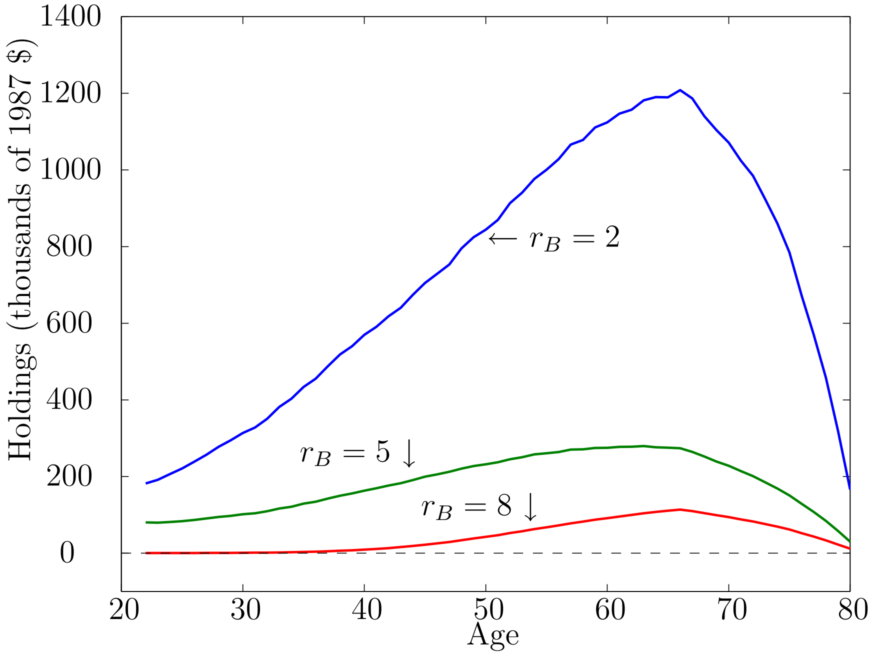

Figure 2: Average equity demand and borrowing as a function of the borrowing rate. Baseline parameter settings.

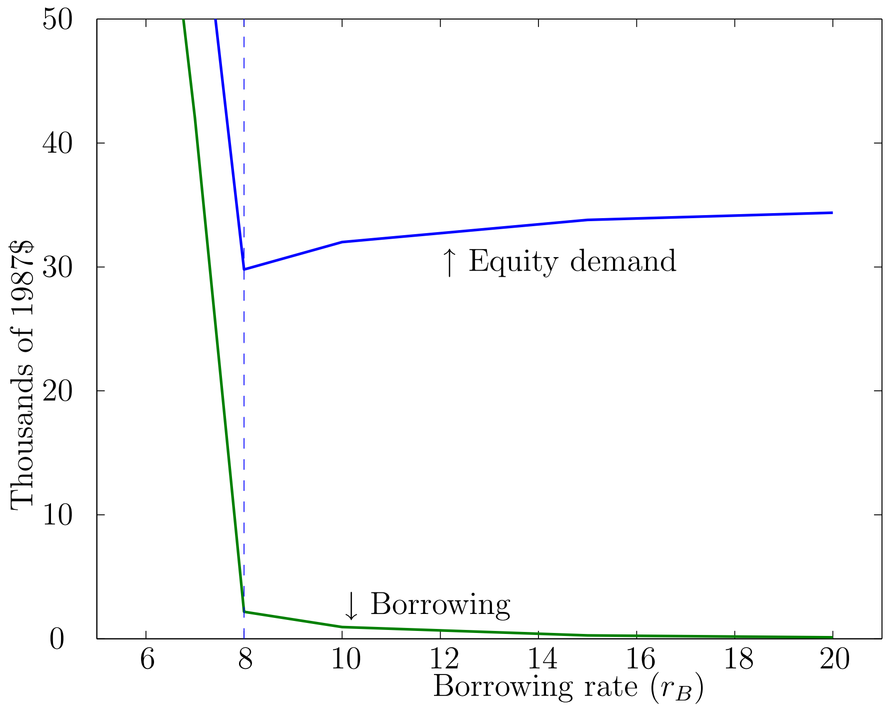

Figure 3: Median equity holdings over the life cycle under alternative borrowing regimes. Baseline parameter settings.

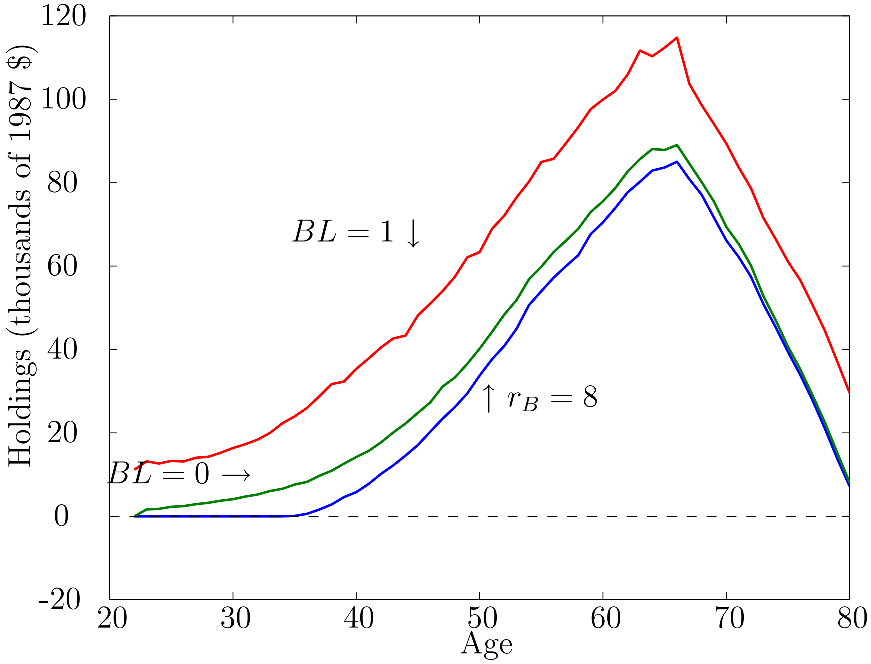

Figure 4: Equity participation rates over the life cycle under alternative borrowing regimes.

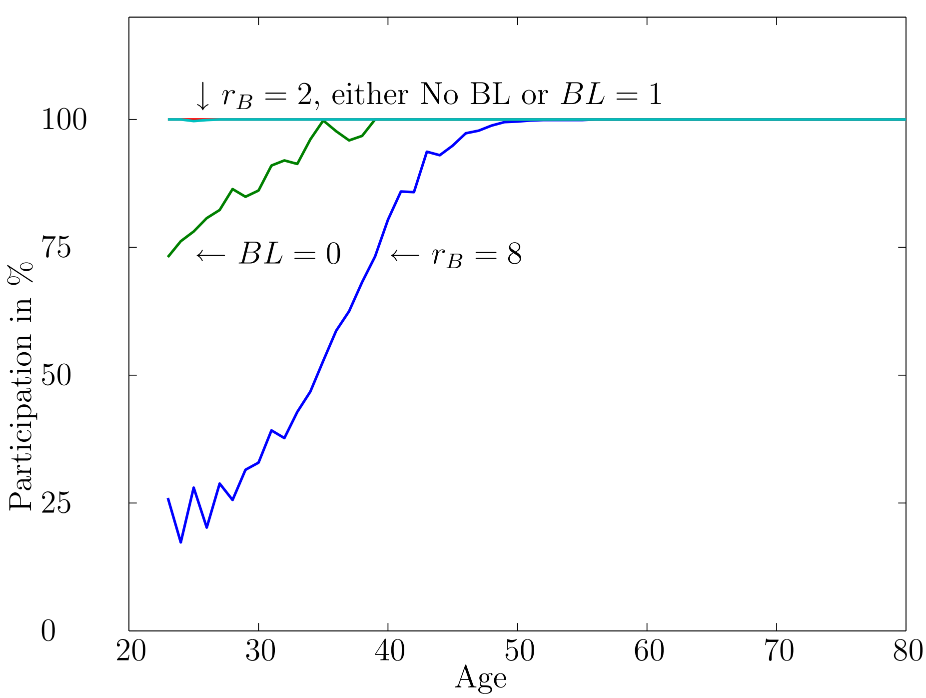

Figure 5: Life-cycle equity holdings for alternative income profiles. $\beta = 0.972$ and no labor income risk as in Specification 3 of Table 4.

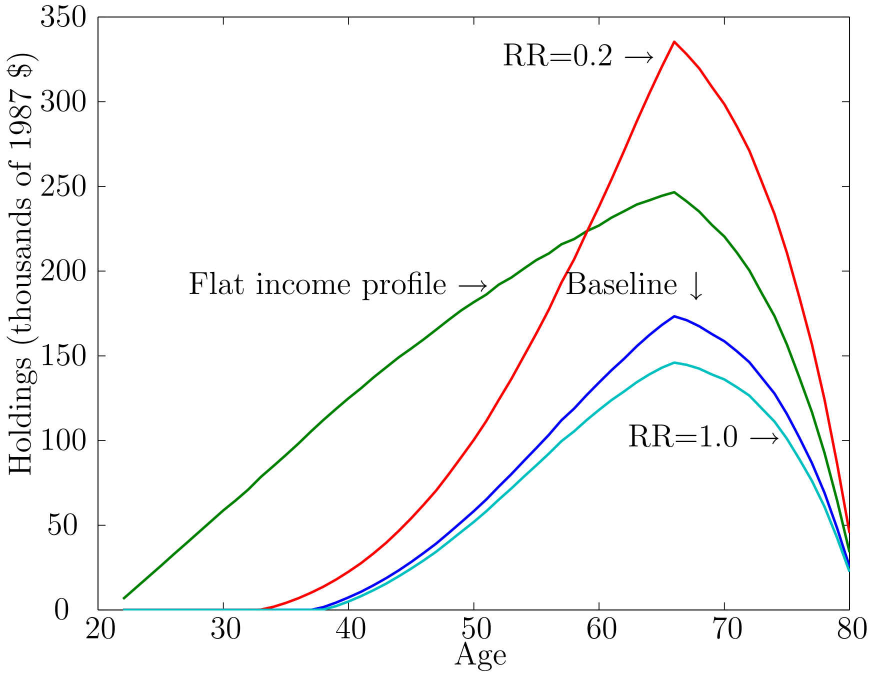

Figure 6: Average equity demand and the variability of permanent income shocks at various borrowing costs. Baseline parameter settings except $\beta = 0.95$.

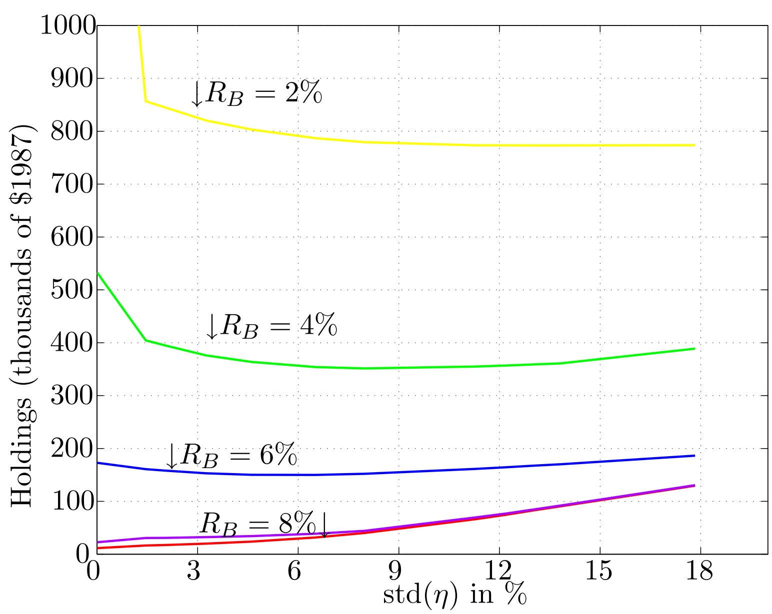

Figure 7: Participation with and without transitory income risk. Baseline setting with no permanent shocks and $\beta = 0.972$.

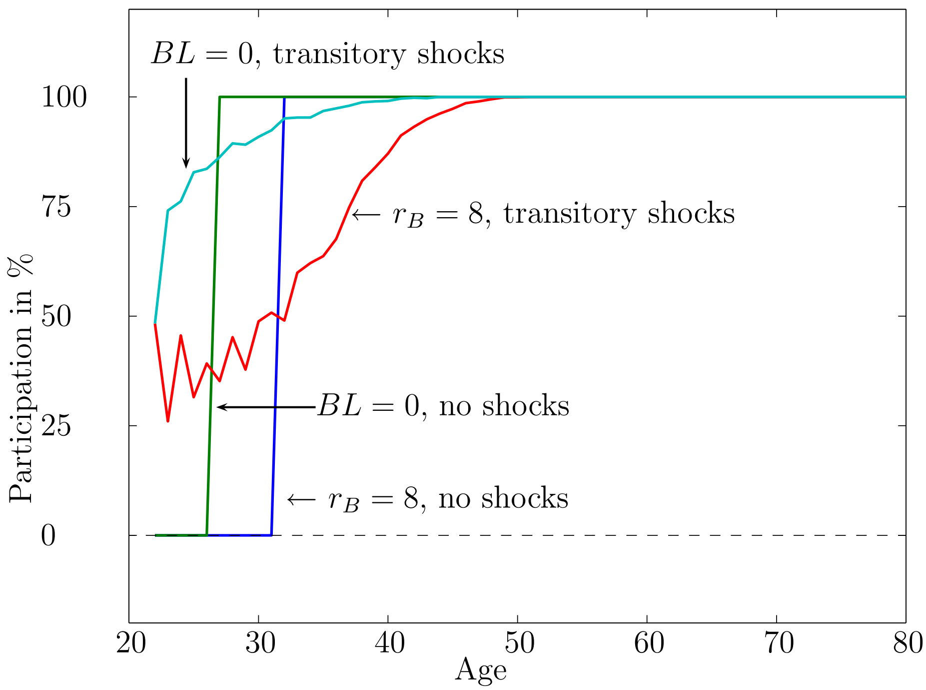

Figure 8: Average equity demand as a function of relative risk aversion. Baseline parameter settings except where noted.