NBER WORKING PAPER SERIES

PORTFOLIO CHOICE OVER THE LIFE-CYCLE IN THE PRESENCE OF 'TRICKLE DOWN' LABOR INCOME

Luca Benzoni
Pierre Collin-Dufresne
Robert S. Goldstein

Working Paper 11247
http://www.nber.org/papers/w11247

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2005

We thank Alexandre Baptista, Marco Cagetti, Graham Candler, Jo˜ao Cocco, Francisco Gomes, John Heaton, Ravi Jagannathan, Ross Levine, Hong Liu, Debbie Lucas, James MacKinnon, Valery Polkovnichenko, Luis Viceira, Amir Yaron, and seminar participants at Washington University, and the University of Virginia for helpful comments and suggestions. We are grateful to Huyan Qiu, who provided excellent research assistance. All errors remain our sole responsibility. The most recent version of this paper can be downloaded from http://www.umn.edu/~ lbenzoni.The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Portfolio Choice over the Life-Cycle in the Presence of ‘Trickle Down’ Labor Income
Luca Benzoni, Pierre Collin-Dufresne, and Robert S. Goldstein
NBER Working Paper No. 11247
March 2005
JEL No. G1, E2, E3

ABSTRACT

Empirical evidence shows that changes in aggregate labor income and stock market returns exhibit only weak correlation at short horizons. As we document below, however, this correlation increases substantially at longer horizons, which provides at least suggestive evidence that stock returns and labor income are cointegrated. In this paper, we investigate the implications of such a cointegrated relation for life-cycle optimal portfolio and consumption decisions of an agent whose non-tradable labor income faces permanent and temporary idiosyncratic shocks. We find that, under economically plausible calibrations, the optimal portfolio choice for the young investor is to take a substantial {\em short} position in the risky portfolio, in spite of the large risk premium associated with it. Intuitively, this occurs because the cointegration effect makes the present value of future labor income flows `stock-like' for the young agent. However, for older agents who have shorter times-to-retirement, the cointegration effect does not have sufficient time to act, and the remaining human capital becomes more `bond-like.' Together, these effects create a hump-shaped optimal portfolio decision for the agent over the life cycle, consistent with empirical observation.

Luca Benzoni
Carlson School of Management
University of Minnesota
321 19th Ave S.
Minneapolis, MN 55455
lbenzoni@umn.edu

Robert S. Goldstein
Carlson School of Management
University of Minnesota
321 19th Ave S.
Minneapolis, MN 55455
rgoldstein@csom.umn.edu

Pierre Collin-Dufresne
Haas School of Business F628
University of California -Berkeley
545 Student Services Bldg #1900
Berkeley, CA 94720-1900
and NBER
dufresne@haas.berkeley.edu

1 Introduction

The optimal portfolio choice problem over the life cycle has received considerable attention from political, financial, and academic circles. In spite of the vast work on this topic, there is still much discord between empirical observation, 'conventional wisdom,' and the predictions of most of the academic literature.

Although the level of stock market participation has increased significantly over the decades, some empirical features have remained robust. In particular, several studies report that risky asset holdings are typically low at young ages, and then are either increasing or hump-shaped as the agent ages (see, e.g., Ameriks and Zeldes (2001), Faig and Shum (2002), Heaton and Lucas (2000), and Poterba and Samwick (2001)). In contrast, conventional wisdom maintains that, for reasonable levels of risk aversion, young agents should place a large proportion of their wealth into the market portfolio, and that this proportion should drop as the agent nears retirement. Indeed, one often-quoted strategy suggested by financial advisors is that investors should place (100 - age)% of their wealth in a well-diversified equity portfolio (see, e.g., Malkiel (1996, p. 418)).

Both empirical observation and conventional wisdom are at odds with early academic studies such as Merton (1969) and Samuelson (1969), who conclude that a long-lived agent should hold a constant fraction of her wealth in the risky asset throughout her life. Moreover, when calibrated using historical values for the equity premium and the stock market return volatility, as well as a 'reasonable' risk-aversion coefficient, these models predict that the appropriate proportion of wealth placed in the risky asset is a counterfactually large number, as high as 100%. These results, however, are derived under many restrictive assumptions, including power utility, independent and identically distributed returns on the risky and riskfree investments, the absence of market frictions and, perhaps most importantly, the absence of labor income.

In an attempt to reconcile theory and observation, many of the restrictive assumptions underlying the Merton (1969) and Samuelson (1969) results have been progressively relaxed.[1] For instance, several studies have examined the effect of labor income on portfolio choice over the life-cycle. For many agents, the 'wealth' (i.e., the certainty-equivalent present value) tied up in terms of future wages dwarfs their financial wealth. As such, one might suspect that optimal portfolio choice ac-

[1] Some papers in this direction examine the implications of time-variation in the riskless interest rate, in the equity premium, and/or different utility functions for the portfolio choice problem. See, e.g., Balduzzi and Lynch (1999), Barberis (2000), Brandt (1999), Brandt et al. (2004), Brennan, Schwartz, and Lagnado (1997), Brennan and Xia (2000), Campbell, Chacko, Rodriguez, and Viceira (2004), Campbell and Viceira (1999, 2001), Dammon, Spatt, and Zhang (2003), Kim and Omberg (1996), Liu (2001), Michaelides (2003), Samuelson (1991), Schroder and Skiadas (1999), Wachter (2002), and Xia (2001).

counting for labor income may generate significantly different predictions. Interestingly, however, most existing studies find that incorporating labor income into the optimal portfolio decision only serves to strengthen the puzzle. Indeed, most models attribute 'bond-like' qualities to the future flow of labor income. That is, these models predict that, through their labor income, agents implicitly hold a large position in the risk-free asset, implying that they should take an even more aggressive position in the risky asset with their cash-on-hand, compared to those models that ignore labor income. Early papers in this direction include Bodie, Merton, and Samuleson (1992), who consider portfolio choice in the context of an endogenous leisure/labor trade-off. More recently, researchers have used micro data to calibrate the individual labor income process. (See, e.g., Campbell et al. (2001, CCGM), Cocco, Gomes, and Maenhout (2002, CGM), Davis and Willen (2000), Haliassos and Michaelides (2003), Jagannathan and Kocherlakota (1996), and Viceira (2001)).[2] With the particular distributional assumptions made in those papers (essentially, labor income and stock returns follow autonomous Markov i.i.d. or AR(1) processes), they find that only counterfactually high correlations between shocks to labor income and stock returns, or the possibility of disastrous labor income shocks (see, e.g., CGM), can explain the low holdings of the risky asset observed for young investors.

We note, however, that the labor income specifications of these models may be unnecessarily restrictive. In particular, if the contemporaneous correlation $\rho_{R_M,L}$ between market returns and changes to aggregate labor income flow is specified to be low, consistent with the data, then these models also force longer-term correlations to be low as well. Closely related, such specifications also force the correlation ρ_{R_M,R_L} between the returns to the market portfolio and returns to human capital (which equals the sum of the current labor income, i.e., 'the dividend', and the unobservable 'capital gain') to be low. In contrast, below we provide evidence suggesting that the correlation between market returns and labor income is an increasing function of the time interval. Note that such a result is consistent with capital flows and flows to labor income being cointegrated, which in turn generates high correlation between the returns to human capital and to the stock market.

The notion that human capital and market returns should be highly correlated is not new. For example, Baxter and Jerman (1997, BJ) test for the existence of cointegration by using data on

[2] Several other contributions investigate the implications of human capital for asset pricing and portfolio choice. For instance, Bodie, Detemple, Otruba, and Walter (2004), Chan and Viceira (2000), and Dybvig and Liu (2004) examine the portfolio choice problem in economies with flexible labor supply or voluntary retirement. Telmer (1993) investigates the variability of the intertemporal marginal rate of substitution in an incomplete market economy with uninsurable labor income shocks. Heaton and Lucas (1996) and Lucas (1994) study the equity premium in economies with aggregate and idiosyncratic labor income shocks, transaction costs, as well as borrowing and short-sales constraints.

aggregate employee compensation and GDP growth (in contrast to using market returns, as is done in this paper). Although the evidence in support of cointegration is statistically weak, as is often the case in tests of cointegration, they proceed under the economically plausible assumption that such a relation indeed exists, and investigate the implications for international portfolio choice. Assuming a constant discount rate, they find that the present values of capital income and labor income exhibit high correlation, in excess of 90%. Campbell (1996) also estimates a similar high correlation between human capital and market returns, but using a very different argument. In particular, he assumes that labor income follows an AR(1) process with low contemporaneous correlation with the stock dividends. However, he assumes that the same (highly time varying) discount factor should be used to discount both labor income and dividends. In his model, the high correlation between human capital and market returns is effectively due to the common highly varying discount factor.[3]

In this paper, we investigate the implication of cointegration between labor income and market returns for life-cycle portfolio choice. Such a specification is consistent with the notion of 'trickle-down' labor income. That is, future labor income flows are affected by past profitability of the economy, so that returns to labor and physical capital are highly correlated, even though the contemporaneous correlations between market returns and changes in labor income might be low.[4]

Although related to the work of BJ, our analysis differs significantly from theirs in many aspects. First, they consider an infinitely-lived representative agent who has a claim to aggregate labor income. Thus, their analysis does not generate implications for the life-cycle behavior of finitely-lived individual agents. Furthermore, their analysis ignores the fact that individual agents face significant idiosyncratic labor income shocks (see, e.g., Carroll and Samwick (1997, CS), CGM, and Gourinchas and Parker (2002, GP)) that are not captured by looking at aggregate averages alone. Second, they do not solve for the optimal portfolio choice. Rather, they focus on the one-period return of an investor desiring a world value weighted (i.e., diversified) portfolio. Finally, they estimate human capital by exogenously setting the discount rate used to discount labor income to a constant.

[3]The assumption that market returns and returns to labor should be highly correlated is also common to many macroeconomic models. While the early paper of Fama and Schwert (1977) had dismissed the empirical relevance of labor income for asset prices (Mayers (1974)), a new strand of literature has recently revisited this point (Black (1995), Jagannathan and Wang (1996), Santos and Veronesi (2004), Campbell (1996), Lettau and Ludvigson (2001)). These recent papers find that labor income can be an important conditioning variable, which improves the predictive power of asset pricing models substantially.

[4]Other papers use the assumption that labor income and dividend flows are cointegrated. For example, Santos and Veronesi (2004) and Menzly, Santos, and Veronesi (2004) investigate economies where the proportion of output paid out as labor income is stationary.

In contrast, we solve the optimal life-cycle portfolio choice problem of an agent with constant relative risk-aversion who earns non-tradable labor income. The latter is cointegrated with stock returns and exhibits both temporary and permanent idiosyncratic labor income shocks, which we model as in CCGM and CGM. Using a dynamic programming approach we solve for the consumption and portfolio allocation rules, and also obtain endogenously the 'shadow' present value of labor income for the optimizing agent (effectively discounting future labor income at her marginal utility).

Contrary to much of the previous literature, and to 'conventional wisdom,' our model predicts that a young agent should take a short position in the risky asset. However, as the agent ages, the optimal proportion of wealth in risky stocks increases. Intuitively, the inverse of the mean reversion coefficient controlling the cointegration provides a time-scale for the agent: if the number of years of remaining employment is larger than this time scale (i.e., if the agent is young), then the return on their human capital is highly exposed to market returns. Furthermore, most of the young agent's 'wealth' is tied up in future labor income. As such, they will find themselves overexposed to market risk, and it will be optimal to short the market portfolio, analogous with the infinitely lived representative agent in BJ who faces no idiosyncratic labor shocks. However, if the number of years of remaining employment is smaller than this time scale (i.e., if they are middle aged), then the return on their human capital is not highly exposed to market returns—that is, their future labor income is more bond-like than stock-like. As such, they find it optimal to invest more in the risky asset than a retired individual. Combined, these results generate a hump-shaped optimal portfolio decision over the life cycle, consistent both qualitatively and quantitatively with empirical evidence.

We emphasize that our results are obtained without specifying any type of fixed entry costs to participate in the equity market. This contrasts considerably with other papers in the literature that can only explain this non-participation by assuming a rather large entry cost (see, e.g., Abel (2001), CGM, and GM). Further, as we demonstrate below, the qualitative conclusions of our findings are very robust across a wide range of parameter inputs. The most important parameter is κ, the mean-reversion coefficient controlling the cointegration between labor income and market returns. Our point estimate for κ ranges from 0.1 to 0.2 depending upon the data set used, and is consistent with the estimates of BJ. Our benchmark case of $\kappa = 0.15$ provides a time-scale of $\frac{1}{0.15} = 6.67$ years for the cointegration to take effect. For times-to-retirement significantly larger than this, the cointegration effect makes the present value of future labor income flows highly correlated with market returns, in turn making it optimal for the young agent to short the risky asset. Interestingly, even when we consider the case $\kappa = 0.05$ (and hence, a time scale on the order of $\frac{1}{0.05} = 20$ years), the same qualitative solution is found for a risk premium of 4% (a number that is often used in the literature).

We acknowledge that we cannot provide irrefutable empirical evidence in support of cointegration (i.e., $\kappa > 0$) over the null hypothesis of a unit root (i.e., $\kappa = 0$). That is, the Dickey-Fuller τ-statistics for unit root tests do not possess the level of significance usually expected in the literature.[5] We note, however, that, as is well known, it is econometrically very difficult to distinguish between these two hypotheses—unit root tests are notorious for lacking power. Still, we consider investigating the implication of such a cointegrated relation for life-cycle portfolio choice a worthwhile endeavor for several reasons. First, we find such a relation economically plausible. As BJ point out, if the labor and capital income were allowed to have independent trends, then the ratio of labor income to capital income would either grow without bound or approach zero asymptotically, and the labor share would approach either zero or one. This seems unlikely (and counterfactual). Second, we provide some additional empirical support for cointegration by demonstrating that correlation between stock returns and labor income are an increasing function of the horizon. Third, we note that our model specification reduces to traditional models (i.e., no cointegration) in the limit $\kappa \to 0$. Econometrically, it is difficult to distinguish between $\kappa = 0$ and, say, $\kappa = 0.05$ given only a few decades of data. Indeed, for $\kappa = 0.05$, we only expect to see the effects of cointegration over a time frame of $\frac{1}{0.05} \approx 20$ years, implying that with 60 years of data, we only have about $\frac{60}{20} \approx 3$ independent data points. Yet, as we show below, the models with $\kappa = 0$ or $\kappa = 0.05$ generate significantly different predictions for the optimal portfolio decision of a young agent.[6] Since both models are difficult to distinguish econometrically, it seems important to investigate the implications of both.

Our conclusions hold for reasonable levels of the agent's risk aversion coefficient. Following CCGM, CGM, and Gomes and Michaelides (2004, GM), we choose $\gamma = 5$ for our baseline case. Qualitatively similar results are obtained if $\gamma = 4$. However, a less risk-averse agent (e.g., $\gamma = 3$) finds it optimal to invest heavily in stocks in spite of the long-run cointegration effect. Hence, we find that even small differences in relative risk aversion can generate substantially different predictions. This result is consistent with empirical observation that asset holdings and stock market participation exhibit a high degree of heterogeneity. In contrast, most models that do not

[5] We note that both our estimates of κ and its statistical significance are very much in line with those obtained by BJ, even though we use stock returns and they use capital flows.

[6] In some respects, this is analogous to the approach of Bansal and Yaron (2004) who show that consumption dynamics with small but persistent drifts are econometrically difficult to distinguish from i.i.d. consumption dynamics, but generate significantly different risk premia.

account for this long-run cointegration conclude that young agents over a wide range of risk-aversion levels should hold a large proportion of their financial wealth in risky securities.

The recent literature has offered many alternative explanations for the limited stock market participation puzzle.[7] The explanation we offer here, while different, can be viewed as complementary to these. Indeed, our paper emphasizes that if, in fact, labor income and market returns are cointegrated over the long run, then such a relation has a first-order effect on the optimal portfolio decisions of an agent over the life cycle.

The rest of the paper is organized as follows. In Section 2, we present the life-cycle portfolio choice model. We explain the details of the model calibration in Section 3. In Section 4 we determine optimal portfolio and consumption choice by numerically solving the Hamilton-Jacobi-Bellman equation. Sensitivity analysis is performed suggesting that the main qualitative result is robust to a wide range of parameter calibrations. We conclude in Section 5.

2 A Model with Cointegrated Aggregate Labor Income

We specify the date-t price of the risky asset as $S(t)$. More accurately, S should be interpreted as the gain process, that is, the value of a portfolio that continually reinvests any dividends paid out by the risky asset. As such, we specify the dynamics of S as if the risky asset does not make dividend payments:

$$\frac{dS(t)}{S(t)} = \mu\, dt + \sigma\, dz_3(t)\,. \tag{1}$$

Here z_3 is a standard Brownian motion. It is convenient to define the log-stock price as $s \equiv \log S$. From Ito's lemma it follows that its dynamics are

$$ds(t) = \left(\mu - \frac{1}{2}\sigma^2\right)\, dt + \sigma\, dz_3(t). \tag{2}$$

Next, we specify the dynamics for the labor income process. Define the current labor income flow for an individual as $L(t)$. It is convenient to introduce log-labor as $\ell(t) = \log L(t)$. Since CCGM, CGM, CS, and GP find idiosyncratic labor shocks to be proportional to the level of labor income, it follows that an individual's income is a product of two numbers: $L_1(t)$, the 'aggregate income' associated with this agent's career choice, and $L_2(t)$, her idiosyncratic shocks. As such, her log-labor flow is a sum of these two factors:

$$\ell(t) = \ell_1(t) + \ell_2(t), \tag{3}$$

[7] See, for example, Abel (2001), Davis, Kubler and Willen (2003), Faig and Shum (2002), GM, Guo (2004), Heaton and Lucas (1997, 2000), Hsu (2003), Hong, Kubik and Stein (2004), and Storesletten et al. (2001).

where $\ell_1(t) = \log L_1(t)$ and $\ell_2(t) = \log L_2(t)$.

We now need to specify the dynamics for $\ell_1(t)$ and $\ell_2(t)$. We choose the process for the aggregate state variable $\ell_1(t)$ to capture two empirical observations: First, contemporaneous correlations between market returns and aggregate shocks to labor income are low. Second, as we report below, this correlation increases substantially with the time horizon.[8]

Thus, we assume that the difference between labor income and the 'gain' process is trend-stationary.[9] (We provide below some empirical evidence that this is indeed consistent with the data.)

Define the difference between the logs of these two variables as $y(t)$:

$$y(t) \equiv \ell_1(t) - s(t). \tag{4}$$

To capture the notion of cointegration (i.e. long-run dependence) between labor income and stock returns, we assume that $y(t)$ is a mean-reverting process with a mean reversion κ and long-run central tendency of $\theta\, t$

$$dy(t) = \kappa\,(\theta\, t - y(t))\; dt + \nu_1\, dz_1(t) - \nu_3\, dz_3(t)\,, \tag{5}$$

where z_1 is a standard Brownian motion independent from z_3.[10]

Note that the time-dependence in the central tendency implies that $\ell_1(t)$ and $s(t)$ are not cointegrated in the strict sense of the word for non-zero θ. However, we emphasize that one should expect such a time trend in the ratio of gain process to income flow (i.e., to find $\theta \neq 0$). To illustrate this point, consider a simple exchange economy where aggregate output is equal to \$1 per year, and the present value of the output is constant and equal to \$20. The price-dividend ratio is stationary (in fact constant) and equal to 1/20=0.05. It follows that the gain process (value of the stock plus reinvested dividends) at date t is simply $20 * (1.05)^t$. The gain process-to-dividend ratio at date-t is $\frac{20*(1.05)^t}{1}$, which is clearly non-stationary. In particular, note that the log gain-dividend ratio

[8] The latter is also consistent with the results of Campbell (1996) and BJ, that the returns to human and physical capital are more highly correlated.

[9] We note that an alternative modeling approach would be to assume that the flow of aggregate labor income and the flow of aggregate dividends are cointegrated, as in, for example, Menzly, Santos, and Veronesi (2004). However, since we are in a partial equilibrium framework, this would force us to specify both the aggregate dividend process and the pricing kernel of the economy in order to determine aggregate stock returns. To avoid this extra layer of structure, rather than specifying a cointegration relation between labor income and dividends,we assume that the difference between labor income and the 'gain' process is 'trend-stationary,' which is also consistent with our empirical results reported below.

[10] It might seem more general to define the y-dynamics as $dy = \kappa\,(\theta_0 + \theta\, t - y)\; dt + \nu_1\, dz_1 - \nu_3\, dz_3$. However, one can then replace the state variable y with $y^* \equiv (y - \theta_0)$, whose dynamics from Ito's lemma follow equation (5) above. Such a transformation of state variables would not affect the definition of y (up to an 'irrelevant constant').

$= \big(\log 20 + t \log 1.05\big)$ is a linear function of time, consistent with the specification of the central tendency in equation (5).

The crucial coefficient in equation (5) is κ, which measures the speed of mean-reversion of the deviations between labor income and stock prices towards a deterministic time trend θt. Roughly speaking, we can think of $\frac{1}{\kappa}$ as a time-scale for which stock returns and labor income are coupled. More formally, $\tau \equiv \log(2)/\kappa$ is the half-life for which deviations from the steady-state value decays. Since a young middle class agent might have, say, 45 years of labor income in front of her before retirement, even a small positive κ implies that the present value of her labor income is strongly affected by market movements, as we demonstrate below.

The AR(1) specification implies that $y(T)$ is normally distributed, with expected value

$$\mathrm{E}_0\left[y(T)\right] = \left[y(0) + \frac{\theta}{\kappa}\right] e^{-\kappa T} + \theta T - \frac{\theta}{\kappa}. \tag{6}$$

This equation shows that the steady state value of $y(t)$ is not θt but rather $\theta t - \frac{\theta}{\kappa}$. Hence, below we choose for our benchmark case $y(0) = -\frac{\theta}{\kappa}$.

Finally, we specify the dynamics of the logarithm of the idiosyncratic shocks to be arithmetic Brownian motion:

$$d\ell_2(t) = \left(\alpha(t) - \frac{\nu_2^2}{2}\right) dt + \nu_2\, dz_{2,i}(t)\,, \tag{7}$$

where $z_{2,i}$ is a standard Brownian motion independent from both z_1 and z_3. The subscript (i) is used to emphasize that this shock is idiosyncratic, in contrast to the aggregate shocks z_1 and z_3. That is, we follow CCGM, CGM, CS, GP, and many others and we assume that the idiosyncratic labor income component is subject to permanent shocks. Further, we introduce a time-dependence in the drift in (7) to capture the findings in the literature that the drift of an individual's labor income is a function of her age. Specifically, we choose

$$\alpha(t) \;\; = \;\; \alpha_0 + \alpha_1 t - \overline{\alpha}\,, \tag{8}$$

where α_0 and α_1 are calibrated to capture the hump-shape of earnings over the life cycle (see, e.g., CGM and CS). However, we note that the combination of $\ell_1(t)$ and $\ell_2(t)$ effectively 'double counts' the expected increase in labor income over time. As such, the term $\overline{\alpha}$ is subtracted so that the total labor income process $\ell(t)$ conforms well to the empirical findings of, e.g., CGM.[11]

From its definition $\ell(t) = y(t) + s(t) + \ell_2(t)$, we find

$$d\ell(t) = \left(\kappa\theta t - \kappa y(t) + \mu - \frac{\sigma^2}{2} + \alpha(t) - \frac{\nu_2^2}{2}\right) dt + \nu_1\, dz_1(t) + \nu_2\, dz_{2,i}(t) + (\sigma - \nu_3)\, dz_3(t). \tag{9}$$

[11] Obviously, only the difference $(\alpha_0 - \overline{\alpha})$ is econometrically identifiable. However, we found it to be convenient for intuitive purposes to distinguish them.

We note that, since the z_1 and $z_{2,i}$ shocks are orthogonal to the stock return shock z_3, the contemporaneous correlation between stock market and labor income shocks is $\mathrm{corr}(ds, d\ell) = \frac{(\sigma-\nu_3)}{\sqrt{\nu_1^2+\nu_2^2+(\sigma-\nu_3)^2}}$. Thus, in the special case $(\sigma - \nu_3) = 0$, labor income is contemporaneously uncorrelated with market returns. We will choose this case as our benchmark case to emphasize that short-term correlations are unnecessary for generating labor income dynamics that are 'stock-like'. Instead, what is crucial is the long-term cointegration.

Equation (9) implies that $\ell(t)$ is normally distributed. Straightforward but tedious algebra gives

$$\mathrm{E}_0\left[\ell(T)\right] = \ell(0) - \left(y(0) + \frac{\theta}{\kappa}\right)\left(1 - e^{-\kappa T}\right) + \left(\theta + \mu - \frac{\sigma^2}{2} + \alpha_0 - \overline{\alpha} - \frac{\nu_2^2}{2}\right) T + \frac{\alpha_1}{2} T^2 \quad (10)$$

$$\mathrm{Var}_0\left[\ell(T)\right] = \left(\frac{\nu_1^2}{2\kappa} + \frac{\nu_2^2}{2\kappa}\right)\left(1 - e^{-2\kappa T}\right) + \left(\sigma^2 + \nu_2^2\right) T - \frac{2\nu_3\sigma}{\kappa}\left(1 - e^{-\kappa T}\right). \quad (11)$$

Normality implies that

$$\begin{aligned} \mathrm{E}_0\left[L(T)\right] &= \mathrm{E}_0\left[e^{\ell(T)}\right] \\ &= e^{\mathrm{E}_0\left[\ell(T)\right] + \frac{1}{2}\mathrm{Var}_0\left[\ell(T)\right]}. \end{aligned} \quad (12)$$

We use this formula to choose $\overline{\alpha}$ to best fit the empirical findings of CGM.

2.1 Comparison with Standard Labor Income Process Specification

Standard approaches typically choose to specify the labor process using levels rather than changes. Furthermore, it is common to consider only discrete time intervals, rather than use a continuous-time specification as above. In order to clarify how our approach relates to the extant literature, here we compare our specifications (9) and (2) for the labor income and stock price with more standard models in the literature. In particular, we demonstrate that in the limit $\kappa \to 0$, our specification is nearly equivalent to the standard model.

For example, CCGM assume that investor's i age t labor income, $Y_{i,t}$, is exogenously given by

$$\log(Y_{i,t}) = f(t, Z_{i,t}) + \nu_{i,t} + \varepsilon_{i,t}\,, \quad (13)$$

where $f(t, Z_{i,t})$ is a deterministic function of age and other individual characteristics $Z_{i,t}$, $\varepsilon_{i,t}$ is an idiosyncratic temporary shock uncorrelated across households and distributed as $N(0, \sigma_\varepsilon^2)$, and $\nu_{i,t}$ is given by

$$\nu_{i,t} = \nu_{i,t-1} + u_{i,t}\,. \quad (14)$$

Here, $u_{i,t}$ is distributed as $N(0, \sigma_u^2)$ and is uncorrelated with $\varepsilon_{i,t}$. Moreover, $u_{i,t}$ is decomposed into an aggregate component ξ_t and an idiosyncratic component $\omega_{i,t}$, uncorrelated across households:

$$u_{i,t} = \xi_t + \omega_{i,t} \,. \tag{15}$$

Further, CCGM assume that the excess return on the risky asset is given by

$$R_{t+1} - \overline{R}_f = \mu + \eta_{t+1} \,, \tag{16}$$

where the innovations η_t are assumed to be i.i.d. over time and distributed as $N(0, \sigma_\eta^2)$. They allow for correlation between the aggregate component of labor income shocks, ξ_t, and innovations to stock returns, η_t; they denote the correlation coefficient $\rho_{\eta,\xi}$.

Using equation (13) at date-t and date-$(t + \Delta t)$, and then using (14), we can write the change in labor income as

$$\begin{aligned}
\log(Y_{i,t+\Delta t}) - \log(Y_{i,t}) &= \left[f(t, Z_{i,t+\Delta t}) - f(t, Z_{i,t})\right] + \left[\nu_{i,t+\Delta t} - \nu_{i,t}\right] + \left[\varepsilon_{i,t+\Delta t} - \varepsilon_{i,t}\right] \\
&= \left[f(t, Z_{i,t+\Delta t}) - f(t, Z_{i,t})\right] + u_{i,t+\Delta t} + \left[\varepsilon_{i,t+\Delta t} - \varepsilon_{i,t}\right] \\
&= \left[f(t, Z_{i,t+\Delta t}) - f(t, Z_{i,t})\right] + \omega_{i,t+\Delta t} + \xi_{t+\Delta t} + \left[\varepsilon_{i,t+\Delta t} - \varepsilon_{i,t}\right] \,. \qquad (17)
\end{aligned}$$

This labor income specification closely matches our specification in equation (9) after some relabeling and some minor changes. Let us ignore for now the temporary shock term $\left[\varepsilon_{i,t+\Delta t} - \varepsilon_{i,t}\right]$. This is done for two reasons. First, we cannot capture this temporary shock in continuous time in the way that CCGM do without introducing another state variable, which would increase significantly the difficulty of obtaining convergence numerically. Instead, we capture the notion of temporary shocks by placing them into the wealth dynamics rather than the labor income dynamics, as will be seen below in equation (22). Second, and more importantly, we emphasize that both CCGM and us find this term to have negligible effect on optimal portfolio decisions. We then relabel $\Delta\ell(t) \equiv \log(Y_{i,t+\Delta t}) - \log(Y_{i,t})$, $\omega_{i,t+\Delta t} \equiv \nu_2\, \Delta z_{2,i}(t)$ and $\left[f(t, Z_{i,t+\Delta t}) - f(t, Z_{i,t})\right] \equiv \left(\mu - \frac{\sigma^2}{2} + \alpha^{(\kappa=0)}(t) - \frac{\nu_2^2}{2}\right)$. Finally, since CCGM allow aggregate labor income shocks ξ to be correlated with innovations in market returns η, we decompose ξ into two terms $\xi_\perp \equiv \nu_1\, \Delta z_1$ and $\xi_\parallel \equiv (\sigma - \nu_3)\, \Delta z_3$, which are 'orthogonal' and 'parallel' to stock market shocks η_t, respectively. Thus, we write $\xi_t \equiv \xi_\perp + \xi_\parallel = \nu_1\, \Delta z_1 + (\sigma - \nu_3)\, \Delta z_3$. With this relabeling and the dropping of the temporary component term, the CCGM and our labor income dynamics can be written, respectively, as

$$\begin{aligned}
\Delta\ell^{CCGM} &= \left(\mu - \frac{\sigma^2}{2} + \alpha^{(\kappa=0)}(t) - \frac{\nu_2^2}{2}\right) \Delta t + \nu_1\, \Delta z_1 + \nu_2\, \Delta z_{2,i} + (\sigma - \nu_3)\, \Delta z_3 \qquad (18) \\
\Delta\ell &= \left(\kappa\,(\theta t - y) + \mu - \frac{\sigma^2}{2} + \alpha^{\kappa}(t) - \frac{\nu_2^2}{2}\right) \Delta t + \nu_1\, \Delta z_1 + \nu_2\, \Delta z_{2,i} + (\sigma - \nu_3)\, \Delta z_3 . \qquad (19)
\end{aligned}$$

Here, the superscript $\alpha^{\kappa}(t)$ emphasizes that we calibrate $\alpha(t)$ for a *given* κ to match the labor income profile of CGM. Clearly, the two models differ only in the conditional drift, and are identical in the limit where the mean reversion parameter $\kappa \to 0$. Below, we demonstrate that even though these two models are extremely difficult to distinguish econometrically for 'small' values of κ, they have enormously different predictions for the optimal portfolio choice of young agents. Indeed, even for an estimate of κ as low as 0.05, which implies a time-scale of $\frac{1}{0.05} = 20$ years, and a risk premium of four percent (the same risk premium assumed by, e.g., CCGM, CGM, and GM), we find it optimal for the young agent to short the market portfolio.

2.2 Empirical Motivation for the Model

BJ use the augmented Dickey-Fuller (ADF) approach to test whether the ratio of labor income to capital income is a stationary random variable. As it is often the case when testing for unit roots, they are unable to find decisive evidence to rule out nonstationarity. As discussed in BJ, these findings can be explained by the fact that unit root tests lack power, i.e., they tend not to reject the unit root null hypothesis when it is false.

We perform similar ADF tests to check whether the variable $y = \ell_1 - s$ is trend stationary. When measuring y, we use return realizations on the U.S. value-weighted market index as a proxy for s, instead of relying on GDP growth data as in BJ. We use average U.S. employee compensation data to construct a proxy for ℓ_1. We obtain the total annual 'Compensation of employees' and the number of 'Full-time and part-time employees' from the National Income and Product Accounts (NIPA) Tables. Each year, we divide the number of 'Full-time and part-time employees' by one minus the average yearly unemployment rate to obtain a proxy for the number of individuals who are either holding or seeking a full- or part-time employee position. Our proxy for ℓ_1 is the logarithm of the ratio of 'Compensation of employees' and the number of individuals either seeking or holding a full- or part-time employee position.[12] Further, in our model the y variable has a time trend. Thus, we follow Oularis et al. (1989) and we include a second-order time trend in the ADF regression model

$$\Delta y = \xi_1 + \xi_2 t + \xi_3 t^2 + \xi_4 y + \Phi(L)\, \Delta y_{-1} + \epsilon\,, \tag{20}$$

which we estimate by ordinary least squares (OLS). Comparing equation (5) and equation (20), we see that if we ignore the effects of the terms $\left(\xi_3 t^2\right)$ and $\left(\Phi(L)\, \Delta y_{-1}\right)$, then ξ_4 would provide

[12] The value-weighted market index returns are from the Center for Research in Security Prices (CRSP) database. We convert stock market prices and aggregate labor income from nominal to real terms by use of the consumer price index (CPI). The NIPA Tables are available from the Bureau of Economic Analysis web site at http://www.bea.doc.gov/bea/dn1.htm, while the CPI index, the unemployment rate, and the civilian participation rate data are available from the St. Louis Federal Reserve Bank web site at http://research.stlouisfed.org/fred2/.

an estimate of $(-\kappa\,\Delta t)$ after one accounts for the transformation from discrete time to continuous time.

Not surprisingly, our analysis suffers from the same problems as those reported in BJ. The ADF test results for the 1948-2001 sample are in Table 2.2 below. The Dickey-Fuller 10% asymptotic

Table 1: ADF test results for the 1948-2001 sample period.

Number of lags in $\Phi(L)$	ξ_4	ADF τ statistic	Centered Adjusted R^2
0	-0.1651	-2.54	0.11
1	-0.1672	-2.39	0.09
2	-0.1520	-2.09	0.14

critical value for the ADF test based on (20) is -3.55 (see, e.g., Davidson and MacKinnon (1993), p. 708). Consistent with BJ (our results are very similar to those that BJ report in their unpublished Appendix), we cannot reject the null hypothesis that the y series has a unit root.

These results are robust to the proxy we use for the aggregate labor income component. For instance, we have also measured ℓ_1 as the logarithm of ratio of 'Compensation of employees' and the number of 'Full-time and part-time employees,' without adjusting the denominator using the unemployment rate. Further, we have projected changes in ℓ_1 against the lagged average yearly job market participation rate of the civilian U.S. population, and we have performed the ADF tests on ℓ_1 minus its predictable component. In each case, we have obtained similar ξ_4 and τ estimates.

Finally, these findings are robust to the sample period. Since unemployment and participation rates are available only from 1948, during the sample period from 1929 to 2001 we can only measure ℓ_1 as the ratio of 'Compensation of employees' and the number of 'Full-time and part-time employees,' without adjusting the denominator using the unemployment rate. This approach yields a $\xi_4 = -0.1858$ estimate with a $\tau = -2.73$ ADF statistic.

To provide additional motivation for the model, we further investigate the properties of the long-run correlations between stock returns and labor income. In Table 2.2, Panel A, we report the correlations between the jth-order differences in log-labor income and log-stock prices, $j = 1, \ldots, 5$, computed using overlapping intervals from 1948 to 2001. First, we note that these correlations are particularly difficult to pin down. For instance, the 1-year correlations range from 16% to 37% depending on the proxy we use for ℓ_1. Further, these estimates vary depending on the choice of the sample period (not reported). This evidence is consistent with the rest of the literature. For instance, CCGM and CGM find that the 1-year correlation is small and insignificant, while they report that the same correlation estimated with returns lagged one year is as large as 52%. Similarly, Davis and Willen find no evidence for a relation between income innovations and contemporaneous aggregate equity returns. Further, Campbell (1996) estimates a VAR model for stock and T-Bill returns, labor income, dividend yields, and term premia. He finds a 17% correlation between the contemporaneous innovations to labor income and stock returns. Clearly, longer-term correlations may be even more imprecisely estimated, as the number of 'independent observations' drops linearly with the time interval for which the correlation is measured.[13] Still, we find two results which seem robust. First, the estimate of the κ coefficient is relatively insensitive to the choice of the proxy for labor income, and is approximately equal to 0.15. Second, and related, we systematically find that correlations are increasing in the time interval. This pattern is at odds with previous labor income models, but is consistent with the notion of cointegration between labor income and stock market performance. To illustrate this point, in Panel B we report the correlations implied by our model (5) and (2) for a realistic calibration of its coefficients (we explain the details of the calibration in Section 3 below). Different rows report long-run correlations for different values of κ and $\mathrm{corr}(d\ell, ds)$, the coefficient of contemporaneous correlation between changes in ℓ and changes in s. For $\kappa \approx 0.15$, the value we use in our baseline calibration, the model produces long-run correlations that are consistent with the empirical evidence. Instead, for $\kappa \to 0$, the model reduces to the standard labor income specifications previously considered in the literature, and produces constant long-run correlations that are at odds with the empirical evidence.

In sum, as in BJ, we cannot provide conclusive evidence to reject the unit root hypothesis in the y series. However, we note that the cointegration effect can act at very low frequencies. For small values of κ, it might take many decades for an agent's wages to catch up with the performance of the economy. Thus, given the relatively short sample period it is not surprising that such effect might go undetected by the ADF test, which is notorious for its lack of power. Further, we emphasize that the stylized evidence on the patterns of the long-run correlations between labor income and stock returns are at odds with the labor income models previously considered in the literature, while they closely match the correlation pattern implied by our model. Finally, we note that economic intuition provides strong support to the notion that labor and capital income are cointegrated. For all these reasons, we proceed under the assumption that the y variable is trend-stationary.

[13] Indeed, this is why we do not report correlations over longer time intervals. For example, with (2001 - 1948 = 53) annual observations, even a 5-year time interval already leaves us with only $\frac{53}{5} \approx 10$ independent observations.

Table 2: Long-run correlations between stock returns and labor income.

	κ	C(1)	C(2)	C(3)	C(4)	C(5)
I	0.13	0.1566	0.3415	0.4126	0.4583	0.4551
II	0.15	0.3713	0.4624	0.5063	0.5904	0.6341
III	0.16	0.1701	0.3878	0.5085	0.6100	0.6438

Panel A: C(j) is the correlation between the jth-order differences in aggregate log-labor income, ℓ_1, and in log-stock prices, s, computed using overlapping intervals. The sample period is 1948-2001. Different rows report correlations for different measures of labor income, which are constructed as follows:

I: ℓ_1 is the logarithm of the average real employee compensation (total number of employees adjusted for the unemployment rate).
II: ℓ_1 is the logarithm of the average real employee compensation minus its component predicted by the lagged job market participation rate among the civilian population.
III: ℓ_1 is the logarithm of the average real employee compensation (total number of employees adjusted for the unemployment rate) minus its component predicted by the lagged job market participation rate among the civilian population.

κ	corr$(d\ell, ds)$	C(1)	C(2)	C(3)	C(4)	C(5)
0.15	0	0.0725	0.1377	0.1955	0.2463	0.2907
0.15	0.13	0.1875	0.2421	0.2898	0.3313	0.3674
0.10	0	0.0491	0.0950	0.1377	0.1770	0.2132
0.10	0.13	0.1677	0.2065	0.2421	0.2746	0.3043
0.05	0	0.0249	0.0491	0.0725	0.0950	0.1168
0.05	0.13	0.1470	0.1677	0.1875	0.2065	0.2247
$\kappa \to 0$	0	0	0	0	0	0
$\kappa \to 0$	0.13	0.13	0.13	0.13	0.13	0.13

Panel B: C(j) is the correlation between the jth-order differences in aggregate log-labor income, ℓ_1, and in log-stock prices, s, implied by our model, equations (5) and (2). Different rows report correlations for different values of the κ coefficient and the contemporaneous correlation between changes in ℓ and changes in s, corr$(d\ell, ds)$. We obtain different values of corr$(d\ell, ds)$ by changing the value of the ν_3 coefficient. The other model coefficients are calibrated to realistic values, as we explain in Section 3 below.

2.3 The agent

The current financial wealth of the agent is stored in two securities. In particular, the agent owns $\theta_0(t)$ shares of the riskless asset, whose date-t price is $B(t)$, and $\theta(t)$ shares of the risky asset, whose date-t price is $S(t)$. Hence, her wealth and wealth dynamics follow

$$\begin{aligned} W(t) &= \theta_0(t)B(t) + \theta(t)S(t) & (21) \\ dW(t) &= -C(t)dt + \theta_0(t)dB(t) + \theta(t)\,dS(t) + L(t)\,dt + \beta W(t)dz_{4,i}(t). & (22) \end{aligned}$$

The last term captures the notion of transient shocks to the agent's wealth. In contrast to most discrete-time models, it is simpler in our continuous-time model to capture these transient shocks in the wealth process rather than in the labor income process. Consistent with intuition, and the numerical results of CGM, we report below that this term has a negligible effect on the agent's consumption and portfolio choices for a wide range of reasonable parameter estimates for β.

It is convenient to define:

$$\begin{aligned} c &= \frac{C}{W} & (23) \\ \pi &= \frac{\theta S}{W} & (24) \\ X &= \frac{L}{W}. & (25) \end{aligned}$$

That is, c is optimal consumption as a percentage of wealth, X represents the ratio of labor income as a percentage of wealth, and π is the optimal proportion of wealth placed in the risky asset. The wealth dynamics can then be written as

$$\frac{dW(t)}{W(t)} = \Big(r + \pi(t)\,(\mu - r) + X(t) - c(t)\Big)\,dt + \pi(t)\sigma\,dz_3(t) + \beta\,dz_{4,i}(t). \quad (26)$$

We assume that the agent has standard constant relative risk aversion utility function. As such, her objective function is:

$$J(t, w(t), L(t), y(t)) \equiv \max_{c,\pi} \mathrm{E}_t \left[\int_t^T du\, e^{-\delta u} \frac{(c(u)e^{w(u)})^{1-\gamma}}{1-\gamma} + \epsilon^{\gamma} e^{-\delta T} \frac{e^{w(T)(1-\gamma)}}{1-\gamma} \right], \quad (27)$$

where $w \equiv \log W$.

The Hamilton-Jacobi-Bellman equation is (dropping time arguments to simplify notation):

$$\begin{aligned} 0 \;=\; & e^{-\delta t}\frac{C^{1-\gamma}}{1-\gamma} + J_t + WJ_W\left(-\frac{C}{W} + r + \pi(\mu - r) + \frac{L}{W}\right) + \frac{1}{2}W^2 J_{WW}\left(\sigma^2\pi^2 + \beta^2\right) \quad (28) \\ & + LJ_L\left(\kappa\theta t - \kappa y + \mu - \frac{\sigma^2}{2} + \alpha(t) - \frac{\nu_2^2}{2} + \frac{1}{2}\left(\nu_1^2 + \nu_2^2 + (\sigma - \nu_3)^2\right)\right) \end{aligned}$$

$$+\frac{1}{2}L^2 J_{LL}\left(\nu_1^2+\nu_2^2+(\sigma-\nu_3)^2\right)+J_y\left(\kappa\theta t-\kappa y\right)+\frac{1}{2}J_{yy}\left(\nu_1^2+\nu_3^2\right)$$
$$+WLJ_{WL}\left(\sigma-\nu_3\right)\pi\sigma-\nu_3\sigma\pi W J_{Wy}+LJ_{Ly}\left(\nu_1^2-\nu_3\left(\sigma-\nu_3\right)\right).$$

The first order conditions for the two controls are

$$0 \;=\; e^{-\delta t}C^{-\gamma}-J_W \tag{29}$$
$$0 \;=\; WJ_W\left(\mu-r\right)+W^2J_{WW}\sigma^2\pi+WLJ_{WL}\sigma\left(\sigma-\nu_3\right)-\sigma\nu_3 W J_{Wy}, \tag{30}$$

leading to the conditions:

$$C \;=\; \left(e^{\delta t}J_W\right)^{-\frac{1}{\gamma}} \tag{31}$$
$$\pi \;=\; -\frac{WJ_W(\mu-r)+WLJ_{WL}\sigma\left(\sigma-\nu_3\right)-\nu_3\sigma W J_{Wy}}{W^2J_{WW}\sigma^2}. \tag{32}$$

Note that equation (31) provides a simple mapping between consumption and J_W. Below, we take advantage of this relation by performing our numerical analysis on partial derivatives on C rather than on J. This is done mostly to improve the stability of our numerical procedure. The added stability can be understood by noting that equation (32) implies that the proportion of wealth placed into the risky asset must be estimated from numerical estimates of the *second* derivative of the value function. Rather, by using equation (31), we can rewrite equation (32) as:

$$\pi=\frac{\frac{C}{W}\left(\mu-r\right)-\frac{L}{W}C_L\sigma\left(\sigma-\nu_3\right)\gamma+\frac{1}{W}\gamma\nu_3\sigma C_y}{\gamma\sigma^2C_W} \tag{33}$$

Note that this relation allows us to determine π by using only first derivatives of C.

As is well known, the CRRA utility function possesses a scaling feature which allows us to eliminate one of the state variables. In particular, for any value of λ, we can write

$$C(\lambda W,\lambda L,y,t)=\lambda\, C(W,L,y,t). \tag{34}$$

Intuitively, this states that if an agent were twice as rich and had twice the labor income, then she would optimally choose to consume twice as much. If we choose $\lambda=\frac{1}{W}$, then we can write

$$C(1,\frac{L}{W},y,t)=\frac{1}{W}C(W,L,y,t). \tag{35}$$

Recall that we have previously defined $X\equiv\frac{L}{W}$ and $c(X=\frac{L}{W},y,t)\equiv C(1,\frac{L}{W},y,t)$. Thus, we can interpret c as the consumption scaled by wealth:

$$c(X=\frac{L}{W},y,t)=\frac{C(W,L,y,t)}{W}. \tag{36}$$

Using standard rules to change variables, we find the optimal portfolio decision can be written in terms of c as

$$\begin{aligned}\pi &= \frac{c(\mu-r)-\gamma\sigma(\sigma-\nu_3)Xc_x+\gamma\nu_3\sigma c_y}{\gamma\sigma^2(c-Xc_X)} \\ &= \frac{\mu-r}{\gamma\sigma^2}+\left(\frac{\mu-r}{\gamma\sigma^2}+\frac{\nu_3}{\sigma}-1\right)\frac{Xc_X}{(c-Xc_X)}+\frac{\nu_3}{\sigma}\frac{c_y}{(c-Xc_X)}. \qquad (37)\end{aligned}$$

The first term is the well known result from Merton (1969). The other terms capture the effects of stochastic labor income and cointegration. Note in particular that for the case where there is no cointegration (i.e., where $\kappa = 0$), it is straightforward to see that the y-state variable is not 'needed,' i.e., $c_y = 0$ and the last term drops out.[14]

Further, we note that it might be difficult for an agent to short-sell securities. Thus, we follow GM, Storesletten et al. (2001), and many others, and we impose the constraint that π lies within the $\pi_{min} = 0$ and $\pi_{max} = 1$ bounds. In Section 4, we relax this constraint and allow the agent to take short positions up to 100% of her financial wealth, i.e., $\pi_{min} = -1$ and $\pi_{max} = 2$.

Just as we have used equation (31) to rewrite the first order condition on π, we can also use this equation to rewrite the Bellman equation. Straightforward but tedious algebra produces the dynamics

$$\begin{aligned}0 = &-c_t+\left(\frac{r-\delta}{\gamma}\right)c-(-c+r+\pi(u-r)+X)(c-Xc_X) \qquad (38)\\ &-\frac{1}{2}\left(\pi^2\sigma^2+\beta^2\right)\left[-(\gamma+1)c^{-1}(c-Xc_X)^2+X^2c_{XX}\right] \\ &-Xc_X\left[\kappa\theta t-\kappa y+\mu-\frac{\sigma^2}{2}+\alpha(t)-\frac{\nu_2^2}{2}+\frac{1}{2}\left(\nu_1^2+\nu_2^2+(\sigma-\nu_3)^2\right)\right] \\ &-\frac{1}{2}X^2\left[-(\gamma+1)c^{-1}c_X^2+c_{XX}\right]\left[\nu_1^2+\nu_2^2+(\sigma-\nu_3)^2\right] \\ &-c_y(\kappa\theta t-\kappa y)-\frac{1}{2}\left(\nu_1^2+\nu_3^2\right)\left[-(\gamma+1)c^{-1}c_y^2+c_{yy}\right] \\ &-\pi\sigma(\sigma-\nu_3)X\left[-(\gamma+1)c^{-1}(c-Xc_X)c_X-Xc_{XX}\right] \\ &+\nu_3\sigma\pi\left[-(\gamma+1)c^{-1}(c-Xc_x)c_y+(c_y-Xc_{Xy})\right] \\ &-\left(\nu_1^2-\nu_3(\sigma-\nu_3)\right)X\left[-(\gamma+1)c^{-1}c_Xc_y+c_{Xy}\right].\end{aligned}$$

[14] Further, note that for the special case where $\mu - r = \gamma\sigma^2$ and $\nu_3 = 0$ we find $\pi(t) = 1$. In that case, the agent invests 100% of her wealth in risky assets irrespective of the correlation between labor income and stock returns, which is driven by ν_1, ν_2! This is a very specific case, where absent any labor income investors would want to invest everything in the stock market (the Merton portfolio is $(\mu - r)/(\gamma\sigma^2) = 1$). With $\nu_3 = 0$ we can think of labor income as giving the agent a random number (determined by ν_1, ν_2) of shares of stock—the agent has no incentive to deviate from his position. Alternatively, we can think in terms of the two effects on the agent's risky asset holding decision that play a role when increasing exposure to labor income risk: background risk vs. diversification motive. They offset each other exactly in this special case.

The 'final condition' is

$$c(X, y, T) = \epsilon^{-1} \qquad \forall (X, y). \tag{39}$$

2.4 Present Value of Labor Income

The first order condition with respect to consumption for the Hamilton-Jacobi-Bellman equation yields

$$J_W = U_C = e^{-\delta t} C^{-\gamma}. \tag{40}$$

Thus, the time-t present value of the agent's labor income is

$$V_t = \mathrm{E}_t \left[\int_t^T ds\, e^{-\delta(s-t)} \left(\frac{C(s)}{C(t)} \right)^{-\gamma} L_s \right]. \tag{41}$$

Below, we estimate the present value to labor income in equation (41) by using the Monte Carlo method.

Given that V_t is a function of only three state variables, namely y, L, and W, we can write the stochastic component of dV as

$$\begin{aligned} dV_{stochastic} &= V_y dy_{stochastic} + V_L dL_{stochastic} + V_W dW_{stochastic} \\ &= \left(\nu_1 V_y + \nu_1 L V_L \right) dz_1 + \nu_2 L V_L\, dz_{2,i} \\ &\quad + \left(-\nu_3 V_y + (\sigma - \nu_3) L V_L + \pi \sigma W V_W \right) dz_3 + \beta W V_W\, dz_{4,i}. \end{aligned} \tag{42}$$

Although there are no traded securities that correlate with the z_1, $z_{2,i}$, and $z_{4,i}$ sources of risk, we can, as a thought experiment, introduce three 'pseudo-securities' X_j, $j = 1, 2$, and 4 such that

$$\begin{aligned} \frac{dX_j(t)}{X_j(t)} &= r\, dt + \sigma dz^Q_{j,i}(t) \\ &= \left(r + \lambda_j(t)\sigma \right) dt + \sigma dz_{j,i}(t), \quad j = 1, 2, 4. \end{aligned} \tag{43}$$

The coefficients $\lambda_j(t)$, $j = 1$, 2, and 4, are the risk premia on these pseudo-securities.[15] We note that if all these claims were traded, then these risk-premia would be pinned down by the observable price processes. In that case markets would be complete and the portfolio problem would have a simple solution (e.g., Duffie (2001)). It is well-known that when markets are incomplete, the incomplete markets portfolio problem can be characterized by a complete markets problem in a fictitiously completed market where the risk premia of the added securities are such that, at the optimum, the agent does not want to hold them (He and Pearson (1991), Karatzas, Lehoczky, Shreve and

[15] For simplicity, we assume these securities pay no dividends and we normalize their diffusion coefficients to be constant (equal to σ). This insures that the securities span all sources of risk.

Xu (1991)). The corresponding risk premia, given the optimal value function, are determined by $\lambda_j(t)dt = -\frac{dX_j(t)}{X_j(t)} \cdot \frac{dJ_W(t)}{J_W(t)}$. Using equation (40), we obtain

$$\lambda_j(t)dt = -\frac{1}{\mathrm{e}^{-\delta t}C(t)^{-\gamma}} d\left(\mathrm{e}^{-\delta t}C(t)^{-\gamma}\right) \cdot dZ_j(t)\,. \tag{44}$$

We then consider a replicating portfolio consisting of an investment θ_S in the stock S, θ_B in the riskfree asset B, and θ_{X_j} in X_j, $j = 1, 2$, and 4:

$$V^{Rep} = \theta_S S + \theta_B B + \theta_{X_1} X_1 + \theta_{X_2} X_2 + \theta_{X_4} X_4\,. \tag{45}$$

The stochastic component of dV^{Rep} is

$$dV^{Rep}_{stochastic} = \theta_S S\sigma dz_3 + \theta_{X_1} X_1 \sigma dz_1 + \theta_{X_2} X_2 \sigma dz_{2,i} + \theta_{X_4} X_4 \sigma dz_{4,i}\,. \tag{46}$$

Thus, by matching coefficients in (42) and (46) we conclude that the proportion of the agent's human capital implicitly tied up in the stock market is

$$\frac{\theta_S S}{V} = \frac{-\nu_3 V_y + (\sigma - \nu_3)LV_L + \pi\sigma W V_W}{\sigma V}\,. \tag{47}$$

Finally, we determine the correlation coefficient between returns to human capital and stock returns, which we denote by ρ. By combining (2) with (42), we obtain

$$\rho = \frac{-\nu_3 V_y + (\sigma - \nu_3)LV_L + \pi\sigma W V_W}{\sigma_V}\,, \tag{48}$$

where $\sigma_V^2 = (\nu_1 V_y + \nu_1 LV_L)^2 + (\nu_2 LV_L)^2 + (-\nu_3 V_y + (\sigma - \nu_3)LV_L + \pi\sigma W V_W)^2 + (\beta W V_W)^2$.

In Section 4, we evaluate (41), (47), and (48) for reasonable model coefficients, and illustrate the effect of 'trickle down' labor income risk on the agent's human capital.

3 Model Calibration

To illustrate the implications of our model, we consider a realistic calibration of its coefficients.

1. *Risky Asset and Riskfree Bond:*

   Consistent with Mehra and Prescott (1985), we fix the real riskfree interest rate at 1% and we assume a 6% risk premium for the risky asset investment, i.e., $r = 1\%$ and $\mu = 7\%$ in real terms. As we will show later, lower estimates of the risk premium make our results stronger in that optimal stock holdings are even lower. Finally, we set $\sigma = 16\%$.

2. *Aggregate Labor Income Dynamics:*

By discretizing (2) and (5), we obtain:

$$\Delta s = a_0 + \xi_1 \tag{49}$$

$$\Delta y = a_1 + a_2 t + a_3 y + \xi_2 , \tag{50}$$

where ξ_1 and ξ_2 are normal error terms and, as defined earlier, $s = \log(S)$ and $y = \ell_1 - s$.

As we explained in more detail in Section 2.2, we consider two sample periods: 1929-2001 and 1948-2001. When measuring y, we use return realizations on the CRSP U.S. value-weighted market index as a proxy for s. For the 1929-2001 sample period, our proxy for ℓ_1 is the logarithm of the ratio of 'Compensation of employees' and the number of 'Full-time and part-time employees,' both from the NIPA Tables. For the 1948-2001 sample period, it is the logarithm of the ratio of 'Compensation of employees' and the number of individuals either seeking or holding a full- or part-time employee position. We convert stock market prices and labor income from nominal to real terms using the CPI index.

We proceed under the assumption that the y variable is trend-stationary. Thus, we first detrend y by estimating the OLS regression $y = \gamma_1 + \gamma_2 t + \varepsilon_y$. Then, we evaluate the residuals $\hat{\varepsilon}_y$ and use OLS to fit the linear model

$$\Delta\hat{\varepsilon}_y = a_3 \hat{\varepsilon}_y + \xi_2 . \tag{51}$$

With this approach, we immediately obtain an estimate for the a_3 coefficient in (50), while a_1 and a_2 are determined by $a_1 = \gamma_2 - a_3\gamma_1$ and $a_2 = -a_3\gamma_2$. Furthermore, we use the time-series of the residuals $\hat{\xi}_1$ and $\hat{\xi}_2$ to estimate the variance terms $\text{var}(\xi_1)$ and $\text{var}(\xi_2)$, as well as the covariance $\text{cov}(\xi_1,\xi_2)$.

Finally, we map the coefficients of the discretized model (49)-(50) into those of the continuous-time diffusions (2) and (5). Specifically, the κ coefficient in (5) is given by $\kappa = -\log(a_3 + 1) = 0.2$ ($\kappa = 0.13$ when we estimate the model using 1948-2001 data). As for the other model coefficients, after making the appropriate conversions we obtain $\theta = -0.0518$, $\nu_1 = 0.0598$, and $\nu_3 = 0.1721$ when using 1929-2001 data ($\theta = -0.0522$, $\nu_1 = 0.0362$, and $\nu_3 = 0.1447$ when using 1948-2001 data).

We use $\kappa = 0.15$ for our baseline case. Interestingly, we document below that even when κ is as small as 0.05 and the market risk premium is fixed at 4% (a value commonly used in the literature), the same qualitative solution is found. Further, we fix $\theta = -0.0518$, $\nu_1 = 0.05$, and $\nu_3 = \sigma = 0.16$. From equations (2) and (9), we see that imposing $\nu_3 = \sigma$ yields a zero contemporaneous correlation between labor income growth and stock market returns. A low correlation is consistent with the empirical evidence reported in, e.g., CGM, Davis and Willen (2000), and Fama and Schwert (1977).

3. *Permanent Idiosyncratic Labor Income Shocks:*

   The variance of the ℓ_2 term is determined by the ν_2 coefficient, which we calibrate to match the magnitude of the typical permanent income component variance, as measured in previous contributions that have modeled the labor income process of individual households by using micro data from the Panel Study of Income Dynamics (PSID). For instance, CGM report values for the standard deviations of the permanent idiosyncratic shocks that range from 0.1 to 0.13, depending on the household's education level. CS's and GP's estimates range from 0.11 to 0.21, depending on the household's occupation and education level. Storesletten et al. (2002) document that the conditional standard deviation of the permanent shocks increases from 0.12 to 0.21 as the economy moves from peak to trough. In our baseline case, we set $\nu_2 = 0.15$. In the next section, we document the sensitivity of our results to different values of ν_2.

   To gather a better sense for the relative magnitude of the ν_2 coefficient in our calibration, it is worth noting that from equation (9) the total variance of the labor income process is $\nu_1^2 + (\sigma - \nu_3)^2 + \nu_2^2$, which equals 0.025 in our baseline case. The total variance can be decomposed into an aggregate and an idiosyncratic component. The aggregate component is $\nu_1^2 + (\sigma - \nu_3)^2 = 0.0025$, while the idiosyncratic component is $\nu_2^2 = 0.0225$. Thus, in our baseline case the ratio of aggregate to permanent idiosyncratic variance shocks is very small, consistent with the evidence in CCGM. In the next section, we document that a larger value of ν_2, which implies an even smaller ratio of aggregate to permanent idiosyncratic variance shocks, also yields qualitatively similar risky asset holdings.

4. *Deterministic Life-Cycle Labor Income Profile:*

   We calibrate the coefficients in the drift term $\alpha(t)$ in (7) to reproduce the typical income pattern due to the predictable growth component described in CS. We consider a twenty year old college-educated agent, $t = 0$, who will work till her retirement date at age 65, $T = 45$. We assume that her $t = 0$ annual labor income is \$15,000 in 1992 USD and we set $\alpha_0 = 0.0722$ and $\alpha_1 = -0.0024$, which imply the deterministic labor income profile depicted in Figure 1. Further, we subtract the term $\overline{\alpha} = 0.0142$ in (8) to compensate for the expected increase in the aggregate labor income component ℓ_1, so that the total labor process $\ell = \ell_1 + \ell_2$ conforms well to the typical deterministic labor income pattern estimated by, e.g., CGM using PSID households data.

5. *Transitory Income Shocks:*

   The transitory income component documented in CGM, CS, GP, and others is built into our model through the term proportional to $dz_{4,i}$ in the wealth dynamics (22). For our baseline case, we fix $\beta = 0.02$, which implies that most transient fluctuations are within $\pm 2\beta$, i.e., $\pm 4\%$, of the current value of wealth. Thus, for an average wealth of, say, \$300,000, most transient shocks will be within $\pm\$12,000$ per year, with a typical yearly shock of $\pm\$6,000$, consistent with the results of, e.g., CGM.

6. *Preferences:*

   The critical parameter in the CRRA utility function is the risk aversion coefficient γ. Mehra and Prescott (1985) argue that reasonable values of γ are smaller than 10. As in CCGM, CGM, and GM, we use $\gamma = 5$ for our baseline case. In the next section, we document the sensitivity of our results to different γ values.

   The magnitude of the remaining coefficients in the value function (27) is less controversial. We follow CCGM, CGM, and GM and we fix $\delta = 0.04$. Cagetti (2003), Dyann, Skinner, and Zeldes (2004), and Hurd (1989) examine the implications of a bequest motive on lifetime saving and consumption decisions. In our application, we follow an approach similar to that of GP and do not explicitly model the agent's behavior during her retirement years. In this case, ϵ determines the number of years of retirement consumption that the investor wants to save for. Accordingly, we calibrate ϵ to generate a wealth accumulation profile over the life cycle that is consistent with the evidence documented in, e.g., Cagetti (2003) for college-educated households. This approach results in $\epsilon = 8$.

7. *Initial Conditions:*

   We consider a twenty year old agent endowed with \$5,000 of cash-on-hand in 1992 USD, i.e., $W(0) = 5$. As mentioned previously, the agent's $t = 0$ annual labor income is \$15,000 in 1992 USD, i.e., $L(0) = 15$. Finally, we fix $y(0)$ at its 'steady state' most likely value, i.e., $y(0) = -\theta/\kappa$, and without loss of generality we initialize the logarithm of the stock market gain process at zero, i.e., $s(0) = 0$.

4 Simulation Results

With the exception of a few special cases,[16] analytic solutions for the life-cycle portfolio choice problem are typically not available. We solve our problem numerically, by using standard finite-difference methods; see, e.g., Ames (1977) and Candler (1999).

Here, we only sketch the numerical solution approach, postponing more details to Appendix A. We solve the consumption problem (38) backwards, starting from the time $T = 45$ terminal condition (39) and going all the way back to the initial date t=0. At each 1/10 of a year, we save the values of c and π on an X- and y-grid. To obtain representative wealth, consumption, investment, and X profiles, we simulate 200,000 W, L, y, and X-paths from their dynamics at the frequency of 1/10 of a year. In the simulations, we fix the controls π and c at the values obtained by interpolating our π and c solutions on the points of the X- and y-grid. Then, we average the realizations of the W, C, π, and X-paths. Finally, we derive analytic solutions for $\mathrm{E}_t[y_s]$ and $\mathrm{E}_t[L_s]$, $t \leq s \leq T$, and use them to determine the representative y and L life-cycle patterns.

4.1 Baseline Case

In Figure 2, Panel A, we report the representative life-cycle wealth, consumption, and labor income profiles that result from our baseline calibration of the model. As expected, accumulated wealth increases over the life of the agent, and her consumption grows proportionally. Finally, the representative individual labor income profile exhibits the typical pattern identified by, e.g., CGM for a college-educated household.

Most interestingly, Panel B of Figure 2 depicts the representative stock holdings, π, over the life-cycle. Contrary to the findings of much of the previous literature, we find that a young agent should not invest in the risky asset. However, as the agent ages, the optimal proportion of wealth in risky stocks increases. Intuitively, the inverse of the mean reversion coefficient controlling the cointegration provides a time-scale for the agent: if the number of years of remaining employment is larger than this time scale (i.e., if the agent is young), then the return on their human capital is highly exposed to market returns. Furthermore, most of the young agent's 'wealth' is tied up in future labor income. As such, she will find herself overexposed to market risk, and it will be optimal to short the market portfolio, analogous with the infinitely lived representative agent in BJ who faces no idiosyncratic labor shocks. Since we impose short-sale constraints, the agent chooses to invest her entire liquid wealth in the riskfree bond. However, if the number of years of remaining

[16] Among recent studies, see, e.g., Duffie et al. (1997), Liu and Loewenstein (2002), Liu et al. (2003), Schroder and Skiadas (2003, 2004).

employment is smaller than this time scale (i.e., if she is middle aged), then the return on her human capital is not highly exposed to market returns—that is, her future labor income is more bond-like than stock-like. As such, she finds it optimal to invest more in the risky asset than a retired individual. Combined, these results generate a hump-shaped optimal portfolio decision over the life cycle, consistent both qualitatively and quantitatively with empirical evidence.

Finally, in Figure 2, Panels A and B, we depict the representative paths for the state variables X and y. It is worth noting the downward sloping X-profile. At a young age, an agent has very limited cash-on-hand relative to her annual labor income (in our calibration, $X = 3$ at $t = 0$). As the agent grows older, her accumulated wealth exceeds her annual labor income and X decreases.

4.2 Human Capital

We use equation (41) to compute the value of a 20-year old agent's human capital. Following the same method discussed previously, we simulate 500,000 wealth and consumption paths and we average across these simulated paths to evaluate (41). For a twenty-year old agent, in the baseline case this approach results in a present value of labor income, V, of approximately \$175,000. Further, we numerically differentiate V with respect to y, L, and W, and use our estimates of V_y, V_L, and V_W to compute the fraction of the agent's human capital tied up in the stock market, as illustrated in (47). We find this fraction to be as large as 54.5%. At first blush, this fraction might not seem high enough to generate our findings, since the optimal retired agent holds about that much in stock, so it would seem that the agent's implicit holdings match her desired holdings, and therefore with her remaining cash-on-hand she should also invest about half of it in the risky asset. However, this estimate does not account for her implicit holdings in the three pseudo-securities X_1, X_2, and X_4 that we introduced in Section 2.4. Figure 4 below shows the decomposition of the replicating portfolio for human capital into its various holdings of stock, pseudo-securities and risk-free money market. We find that the position in X_1, X_2, and X_4 implicit in the agent's human capital are 14.4%, 94.2%, and 0.4%, respectively. Clearly, human capital is mostly equivalent to a long position in the stock market S and in permanent idiosyncratic risk which is hedged with X_2. The transient idiosyncratic shocks driven by z_4 and hedged with X_4 represents only a very small fraction of the replicating portfolio. Hence, they do not affect much the shadow value of labor income. We emphasize that the pseudo-securities have risk-premia determined endogenously, so that agents, given their labor income, do not want to trade in these securities.[17] Interestingly, through her

[17] An alternative interpretation is the following: suppose the agent had no labor income, but instead could invest in these pseudo-securities (with risk-premia as determined above), then she would want to invest precisely in the portfolio represented in Figure 4.

human capital the agent's implicit holding in the risk-free asset is approximately -63%. That is, the agent's present value of labor income is a very leveraged security. On the other hand, for an agent approaching retirement human capital becomes small. Thus, her position in these pseudo securities approaches zero, which explains her long position in the stock market.

Related, we measure the correlation of stock returns and the returns to human capital. Using equation (48), for a twenty-year old agent we find a correlation coefficient $\rho \cong 50\%$. That is, due to the idiosyncratic labor income shocks, the correlation is much lower than what is found by BJ and Campbell (1996) at the aggregate level. Still, it is sufficiently high to have a first-order effect on the agent's portfolio choice decisions.

Finally, in Figure 5, Panel A, we illustrate how the agent's human capital evolves over the life cycle. For values of time t from 0 to 45, we use (41) to compute the present value of the future stream of labor income, V_t. We note that the fraction of the agent's labor income tied up in the risky asset is roughly constant at 50% throughout the first half of her life, and it rapidly goes to zero as she approaches retirement. Further, we note that the present value of human capital has a hump-shaped profile. That is, although young agents face a larger stream of future labor income, they discount such cash flows higher than older agents do. This occurs for two reasons. First, as the agent ages, she faces lower idiosyncratic labor income risk. To validate this intuition, we use equation (44) to compute the risk premium on the permanent idiosyncratic labor income shocks over the agent's life cycle. Figure 5, Panel B, shows that λ_2 has a downward sloping profile and confirms that the risk premium on idiosyncratic shocks approaches zero when the agent retires. This effect is common to other models with idiosyncratic labor income risk, e.g., CCGM, CGM, CS, GM, and more. Second, in our model human capital has pronounced 'stock-like' features, and thus it commands a higher discount rate, for young agent, while it acquires 'bond-like' properties, and thus it is discounted at a lower rate, for older agents. Due to this second effect, which is determined by the long-run cointegration of labor income and stock market performance, the value of human capital peaks at a later point of the agent's life, compared to standard models considered in previous studies. This intuition is confirmed by the evidence in Figure 5, Panel C, which shows that the correlation of stock returns and the returns to human capital remains high and basically constant over the first half of the agent's life, and it rapidly drops as the agent approaches retirement.

4.3 Speed of Mean Reversion and Equity Premium

In Figure 6, we explore the robustness of our results to the magnitude of the κ coefficient. Consistent with the intuition discussed in Section 4.1, we see that larger values of κ increase the agent's exposure to stock market risk and thus reduce her stock holdings. However, even a small value of κ has first-order effects on the life-cycle π profile.

CCGM, CGM, and GM set the equity premium equal to 4%, a value that can be motivated based on the observation that stock prices have tended to increase over recent years relative to corporate earnings. Thus, in Figure 7 we illustrate the life-cycle π profile when $r = 2\%$ and $\mu = 6\%$. Interestingly, a lower value of the equity premium makes our results even stronger. Specifically, it is worth noting that with this model calibration a young agent chooses not to invest in the stock market even if the κ coefficient is as low as 0.05, as compared to the $\kappa = 0.15$ of the baseline case.

4.4 Contemporaneous Correlation of Stock Returns and Aggregate Labor Income Shocks

We noted previously that our baseline calibration implies a zero contemporaneous correlation of stock returns and growth rates in labor income. In Figure 8, we illustrate the effect of non-zero contemporaneous correlations. We consider two cases, $\nu_3 = 0.18$ and 0.14, which imply correlations of approximately $\pm 13\%$, respectively. Consistent with previous studies, we note that even such high values of correlations have limited impact on the agent's stock holdings, compared to the long-run cointegration effect.

4.5 Persistent Idiosyncratic Labor Income Shocks

We note that an increase in the idiosyncratic labor income variance (through an increase in ν_2) has two possibly opposite effects on the investor's desired portfolio holdings. First, it increases 'background risk,' which all else equal leads to a decrease in desired risky asset holdings. Second, it provides a 'diversification motive,' which might induce the agent to increase her demand of the risky asset. The latter effect could potentially counter-balance the effect due to the long-run cointegration-like behavior of the aggregate labor income with the market portfolio. In Figure 9, we show that a value of ν_2 as high as 0.20 (the upper end of the empirical range documented in the literature) attenuates but does not eliminate our main result. Interestingly, the picture shows that investors with an investment horizon of approximately 12 years are in fact indifferent to a change in ν_2. This 'duration' like feature may be due to a near perfect offsetting of the two effects (diversification motive vs. background risk) noted above.

4.6 Transitory Idiosyncratic Labor Income Shocks

It is generally agreed that transient idiosyncratic labor income shocks have negligible implications for the optimal portfolio choice problem solution. In Figure 10, we confirm this result by considering values of β as small as zero, and as large as 0.04 (twice the value used in our baseline case).

4.7 Relative Risk Aversion

In Figure 11, we document the sensitivity of our results to changes in the relative risk aversion coefficient. We note that even for a young agent with relative risk aversion $\gamma = 4$ human capital has stock-like features. In this case, the stock holdings retain the same hump-shaped profile over the life-cycle. However, a less risk averse agent (e.g., $\gamma = 3$) perceives her human capital to be more bond-like, in spite of the long-run cointegration effect. Even at a young age she invests heavily in the risky asset. As she gets older, the present value of her human capital declines relative to the value of her liquid wealth. Thus, we see her π profile decline as she approaches retirement. We emphasize that even small differences in relative risk aversion can generate opposite decision as to whether to participate in the stock market. This result is qualitatively consistent with empirical evidence, which shows that asset holdings exhibit a high degree of heterogeneity.

4.8 Short-Sale Constraints

The recent development of derivatives markets as well as the proliferation of Exchange Traded Funds (ETFs) makes it easier for an agent to take short positions in the market portfolio. Thus, in Figure 12 we illustrate the typical life-cycle investment profile when the short-sale constraint is relaxed. Consistent with the intuition developed in Sections 4.1 and 4.2, we find that a young agent chooses to short the market portfolio, to hedge the long position in the stock market implicit in her human capital.

5 Conclusions

Conventional wisdom maintains that young investors should invest heavily in the stock market. Most theoretical investigations concur. Furthermore, most models suggest that labor income is more 'bond-like' than 'stock-like,' implying even higher optimal proportions of wealth should be placed into the risky asset if labor income is accounted for. In this paper, however, we claim that by incorporating two features previously documented in the literature, namely, that aggregate labor income is cointegrated with aggregate output, and that individual labor income is subject to significant permanent idiosyncratic shocks, we find the optimal portfolio choice for the young investor is to take a substantial short position in the risky portfolio. This occurs because our model implies that the value of the claim to labor income is effectively a highly leveraged security with large implicit exposure to the market portfolio.

One obvious extension of our paper is to include housing.[18] We note that Quan and Titman (1997) find that the real estate market is cointegrated with the stock market. This evidence suggests that if one were to incorporate housing into the portfolio choice and model this cointegration, the optimal investment in stocks would become even more negative.

Although this paper focuses on the individual's optimal portfolio and consumption choices given the risk premium of the market as given, our findings might have important implications for general equilibrium models that attempt to explain the 'equity premium puzzle'.[19] Indeed, as pointed out by Basak and Cuoco (1998), by taking as given that a large proportion of investors do not participate in the stock market, one need only attribute very reasonable levels of risk aversion to those agents that do invest in stocks in order to explain the historical equity premium. Our results indicate that it is optimal for a large proportion of agents in the economy to short, or at least not participate in the equity market. Thus, the exogenous specification of Basak and Cuoco (1998) might be justifiable in a general equilibrium setting that considers two agent classes that endogenously choose to participate in the stock market depending on their risk aversion and long-run exposures to aggregate risk.[20]

Further, since we find that, in the presence of cointegration, the investment horizon has a dramatic impact on portfolio holdings, this suggests it would be interesting to understand, within an equilibrium model, the interaction of various cohorts or overlapping generations of households whose labor income is cointegrated with long-term market performance.[21] Within this setting, it would be interesting to examine the effect of possible changes to the Social Security system, e.g., the possibility of moving to a privatized retirement system in which retirement contributions earn market-based rates of return (see, e.g., Abel (2001) and CCGM).

Finally, our model suggests that labor income artificially generates a negative net supply of risk-free securities. This prediction contrasts with the typical approach of assuming that the risk free security is in zero net supply. We save these interesting questions for future research.

[18] Several recent studies investigate the implications of real estate holding for asset pricing. See, e.g., Cocco (2000), Flavin and Yamashita (2002), Hu (2002), Davidoff (2003), and Yao and Zhang (2003).

[19] Related work includes Mankiw and Zeldes (1991), Polkovnichenko (2004), Vissing-Jørgensen (2002), Guvenen (2004) and Vissing-Jørgensen and Attanasio (2003). Although not directly related, Bansal and Yaron (2004) investigate how long run AR(1) processes for earnings flow can explain historical equity premiums.

[20] More specifically, the fraction of human capital implicitly tied up in the stock market might vary by occupation. This effect, which is captured by different values of the κ coefficient in our model, has a significant impact on portfolio holdings. Further, we have demonstrated above that small differences in risk aversion can also yield heterogeneity in stock market participation in our model.

[21] For related work, see, e.g., Constantinides, Donaldson, and Mehra (2002), Guvenen (2004), and Storesletten, Telmer and Yaron (2003).

Appendix A: Numerical Solution Approach

We solve the optimal portfolio and consumption problem (37), (38), and (39) by using the alternate direction implicit (ADI) finite-difference method; see, e.g., Ames (1977). We follow Candler (1999) and treat the non-linear terms in (38) 'explicitly,' thus reducing the problem to a sequence of tri-diagonal systems of linear equations that can be solved easily using standard numerical methods.[22]

As noted previously, via some transformations we are able to reduce the state space from four state variables to two, namely, X and y. We evaluate the solution on a discrete state-space grid. For y, we set the lower bound of the domain at $y_{min} = y_0 + \theta\, T - 3\,\sigma(y)$, and the upper bound at $y_{max} = y_0 + 3\,\sigma(y)$, where $\sigma(y) = \sqrt{(\nu_1^2 + \nu_3^2)/2\,\kappa}$. We then construct the y-grid with a $\Delta y = 0.05$ mesh. For X, we use $X_{min} = 0$ and $X_{max} = 10$, and construct the corresponding X-grid using a $\Delta X = 0.05$ mesh.

We solve the problem backwards, starting from the time $T = 45$ terminal condition (39) and going all the way back to the initial date $t = 0$. We use a time step $\Delta t = 0.0005$, which is further broken down into time-increments of length $\Delta t/2$ in each of the two steps of the ADI algorithm.

We note that our numerical approach is robust to the choice of the time- and space-grid parameters. For instance, we have verified that using a finer $\Delta X = \Delta y = 0.01$ mesh, in combination with different values of X_{max}, y_{min}, and y_{max}, results in the same numerical solution for c and π.

The boundary conditions are treated as follows. First, we note that at $X_{min} = 0$ labor income is zero. Thus, the Merton (1969) closed-form solution for optimal consumption holds and provides an exact boundary condition, which we impose in our finite difference approach. Further, we note that the second derivative of consumption with respect to the X state variable vanishes as X increases. Thus, we impose the condition

$$\frac{\partial^2 c(X_{max}, y)}{\partial X^2} = 0. \tag{52}$$

Economic intuition does not offer exact boundary conditions at y_{min} and y_{max}. After some experimentation, we have found that the third derivative of consumption with respect to the y variable vanishes as y approaches the boundaries of its domain. Thus, we impose the conditions

$$\frac{\partial^3 c(X, y)}{\partial y^3} = 0\,,\; y = y_{min} \text{ and } y = y_{max}. \tag{53}$$

We check the robustness of the solution to this approach by extending the range of the y-domain, finding identical results. Further, we note that using a discretization of (38) that relies only on

[22] We test our numerical approach in the special case of the Merton's (1969) model, for which a closed-form solution is known. In that case, the approximation error generated by the numerical solution method for the agent's consumption/investment policies is nearly zero.

internal points at y_{min} and y_{max} yields results identical to those obtained by imposing the boundary condition (53).

6 References

Abel, Andrew B., 2001, The Effects of Investing Social Security Funds in the Stock Market When Fixed Costs Prevent Some Households from Holding Stocks, *American Economic Review* 91, 128-148.

Ameriks, John, and Stephen P. Zeldes, 2001, How Do Household Portfolio Shares Vary with Age?, Working paper, Columbia University.

Ames, William F., 1977, Numerical Methods for Partial Differential Equations, Academic Press, London.

Balduzzi, Pierluigi, and Anthony W. Lynch, 1999, Transaction Costs and Predictability: Some Utility Cost Calculations, *Journal of Financial Economics* 52, 47-78.

Bansal, Ravi, and Amir Yaron, 2004, Risks For The Long Run: A Potential Resolution of Asset Pricing Puzzles, *Journal of Finance*, 59, 1481-1510.

Barberis, Nicholas, 2000, Investing for the Long Run when Returns Are Predictable, *Journal of Finance* 55, 225-264.

Basak, Suleyman, and Domenico Cuoco, 1998, An Equilibrium Model with Restricted Stock Market Participation, *Review of Financial Studies* 11, 309-341.

Baxter, Marianne, and Urban J. Jermann, 1997, The International Diversification Puzzle Is Worse Than You Think, *American Economic Review* 87, 170-180.

Black, Fisher, 1995, *Exploring General Equilibrium*, MIT Press, Cambridge, Mass..

Bodie, Zvi, Robert C. Merton, and William F. Samuelson, 1992, Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model, *Journal of Economic Dynamics and Control* 16, 427-449.

Bodie, Zvi, Jérôme B. Detemple, Susanne Otruba, and Stephan Walter, 2004, Optimal ConsumptionPortfolio Choices and Retirement Planning, *Journal of Economic Dynamics and Control* 28, 1115-1148.

Brandt, Michael W., 1999, Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach *Journal of Finance* 54, 1609-1645.

Brandt, Michael W., Amit Goyal, Pedro Santa-Clara, and Jonathan R. Stroud, 2004, A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Return Predictability *Review of Financial Studies*, forthcoming.

Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic Asset Allocation, *Journal of Economic Dynamics and Control* 21, 1377-1403.

Brennan, Michael J., and Yihong Xia, 2000, Stochastic Interest Rates and the Bond-Stock Mix, *European Finance Review* 4, 197210.

Cagetti, Marco, 2003, Wealth Accumulation over the Life Cycle and Precautionary Savings, *Journal of Business and Economic Statistics* 21, 339-353.

Campbell, John Y., 1996, Understanding Risk and Return, *Journal of Political Economy* 104, 298-345.

Campbell, John Y., George Chacko, Jorge Rodriguez, and Luis M. Viceira, 2004, Strategic Asset Allocation in a Continuous-Time VAR Model, *Journal of Economic Dynamics and Control*, forthcoming.

Campbell, John Y., João Cocco, Francisco J. Gomes, and Pascal Maenhout, 2001, Investing Retirement Wealth: A Life-Cycle Model, in John Y. Campbell and Martin Feldstein, eds.: Risk Aspects of Investment-Based Social Security Reform, The University of Chicago Press, Chicago, Ill..

Campbell, John Y., and Luis M. Viceira, 1999, Consumption and Portfolio Decisions When Expected Returns are Time Varying, *Quarterly Journal of Economics* 114, 433-495.

Campbell, John Y., and Luis M. Viceira, 2001, Who Should Buy Long-Term Bonds?, *American Economic Review* 91, 99-127.

Candler, Graham V., 1999, Finite-Difference Methods for Continuous-Time Dynamic Programming, in Ramon Marimon and Andrew Scott, eds.: Computational Methods for the Study of Dynamic Economies, Oxford University Press, New York.

Carroll, Christopher D., and Andrew A. Samwick, 1997, The Nature of Precautionary Wealth, *Journal of Monetary Economics* 40, 41-71.

Chan, Yeung Lewis, and Luis M. Viceira, 2000, Asset Allocation with Endogenous Labor Income: The Case of Incomplete Markets, Working Paper, Harvard University.

Cocco, João, 2000, Hedging House Price Risk with Incomplete Markets, Working Paper, London Business School.

Cocco, João, Francisco J. Gomes, and Pascal J. Maenhout, 2002, Consumption and Portfolio Choice over the Life-Cycle, *Review of Financial Studies*, forthcoming.

Constantinides, George M., John B. Donaldson, and Rajnish Mehra, 2002, Junior Can't Borrow: A New Perspective on the Equity Premium Puzzle, *Quarterly Journal of Economics* 117, 269-296.

Constantinides, George M., and Darrell Duffie, 1996, Asset Pricing with Heterogeneous Consumers, *Journal of Political Economy* 104, 219-240.

Dammon, Robert M., Chester S. Spatt, and Harold H. Zhang, 2003, Capital Gains Taxes and Portfolio Rebalancing, *Research Dialogue of TIAA-CREF Institute* 75, 1-14.

Davidoff, Thomas, 2003, Labor Income, Housing Prices and Homeownership, Working Paper, Haas School of Business.

Davidson, Russell, and James G. MacKinnon, 1993, Estimation and Inference in Econometrics, Oxford University Press, New York.

Davis, Steven J., and Paul Willen, 2000, Occupation-Level Income Shocks and Asset Returns: Covariance and Implications for Portfolio Choice, CRSP Working Paper number 523 and NBER Working Paper number 7905.

Duffie, Darrell, 2001, *Dynamic Asset Pricing Theory*, Princeton University Press, Princeton and Oxford.

Duffie, Darrell, and Rui Kan, 1996, A Yield-Factor Model of Interest Rates, *Mathematical Finance* 6, 379-406.

Duffie, Darrell, Wendell Fleming, Mete Soner, and Thaleia Zariphopoulou, 1997, Hedging in Incomplete Markets with HARA Utility, *Journal of Economic Dynamics and Control* 21, 753-782.

Dynan, Karen, Jonathan Skinner, and Stephen P. Zeldes, 2004, Do the Rich Save More? *Journal of Political Economy* 112, 397-444.

Dybvig, Philip H., and Hong Liu, 2004, Lifetime Consumption and Investment: Retirement and Constrained Borrowing, Working Paper, Washington University in Saint Louis.

Faig, Miquel, and Pauline Shum, 2002, Portfolio Choice in the Presence of Personal Illiquid Projects *Journal of Finance* 57, 303-328.

Fama, Eugene F., and G. William Schwert, 1977, Human Capital and Capital Market Equilibrium, *Journal of Financial Economics* 4, 95-125.

Flavin, Marjorie, and Takashi Yamashita, 2002, Owner-Occupied Housing and the Composition of the Household Portfolio, *American Economic Review* 92, 345-362.

Gomes, Francisco J., and Alexander Michaelides, 2004, Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence, *Journal of Finance*, forthcoming.

Gourinchas, Pierre-Olivier, and Jonathan A. Parker, 2002, Consumption Over the Life Cycle, *Econometrica* 70, 47-89.

Guo, Hui, 2004, Limited Stock Market Participation and Asset Prices in a Dynamic Economy, *Journal of Financial and Quantitative Analysis* 39, 495-516.

Guvenen, Fatih, 2004, A Parsimonious Macroeconomic Model for Asset Pricing: Habit Formation or Cross-Sectional Heterogeneity?, Working Paper, University of Rochester.

Haliassos, Michael, and Alexander Michaelides, 2003, Portfolio Choice and Liquidity Constraints, *International Economic Review* 44, 143-177.

Heaton, John, and Deborah J. Lucas, 1996, Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing, *Journal of Political Economy* 104, 443-487.

Heaton, John, and Deborah Lucas, 2000, Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk, *Journal of Finance* 55, 1163-1198.

Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein, 2004, Social Interaction and Stock-Market Participation, *Journal of Finance* 59, 137-163.

Hsu, Jason C., 2003, What Drives Equity Market Non-participation?, Working Paper, UCLA.

Hu, Xiaoqing, 2002, Portfolio Choice for Home Owners, Working Paper, Northwestern University.

He, Hua, and Neil Pearson, 1991, Consumption and Portfolio Policies with Incomplete Markets and Short-sale Constraints: The Infinite Dimensional Case, *Journal of Economic Theory* 54, 259304.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, 1995, Precautionary Saving and Social Insurance, *Journal of Political Economy* 103, 360-399.

Hurd, Michael D., 1989, Mortality Risk and Bequests, *Econometrica* 57, 779-813.

Jagannathan, Ravi, and Narayana R. Kocherlakota, 1996, Why Should Older People Invest Less in Stocks than Younger People?, *Federal Reserve Bank of Minneapolis Quarterly Review* 20, 11-23.

Jagannathan, Ravi, and Zhenyu Wang, 1996, The Conditional CAPM and the Cross-Section of Expected Returns, *Journal of Finance* 51, 3-53.

Karatzas, Ioannis, John P. Lehoczky, Steven E. Shreve, and Gan-Lin Xu, 1991, Martingale and duality methods for utility maximization in an incomplete market, *SIAM J. Control Optim.* 29-3, 702-730.

Kim, Tong Suk, and Edward Omberg, 1996, Dynamic Nonmyopic Portfolio Behavior, *Review of Financial Studies* 9, 141-161.

Lettau, Martin, and Sydney Ludvigson, 2001, Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying *Journal of Political Economy* 109, 1238-1287.

Liu, Hong and Mark Loewenstein, 2002, Optimal Portfolio Selection with Transaction Costs and Finite Horizons, *Review of Financial Studies* 15 , 805-835.

Liu, Jun, 2001, Portfolio Selection in Stochastic Environments, Working Paper, UCLA.

Liu, Jun, Francis A. Longstaff, and Jun Pan, 2003, Dynamic Asset Allocation with Event Risk," *Journal of Finance* 58, 231-259.

Lucas, Deborah J., 1994, Asset Pricing with Undiversifiable Income Risk and Short Sales Constraints: Deepening the Equity Premium Puzzle, *Journal of Monetary Economics* 34, 325-341.

Malkiel, Burton G., 1996, A Random Walk Down Wall Street, Including a Life-Cycle Guide to Personal Investing, W. W. Norton & Company, New York.

Mankiw, N. Gregory, and Stephen P. Zeldes, 1991, The Consumption of Stockholders and Non-stockholders, *Journal of Financial Economics* 29, 97-112.

Mayers, David, 1974, Portfolio Theory, Job Choice and the Equilibrium Structure of Expected Wages, *Journal of Financial Economics* 1, 23-42.

Mehra, Rajnish, and Edward C. Prescott, 1985, The Equity Premium: A Puzzle, *Journal of Monetary Economics* 15, 145-161.

Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding Predictability, *Journal of Political Economy* 112, 1-47.

Merton, Robert C., 1969, Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case, *Review of Economics and Statistics* 51, 247-257.

Michaelides, Alexander, 2003, Portfolio Choice, Liquidity Constraints and Stock Market Mean Reversion, Working Paper, London School of Economics.

Ouliaris, Sam, Peter C. B. Phillips, and Joon Y. Park, 1989, Testing for a Unit Root in the Presence of a Maintained Trend, in Raj, Baldev, ed., Advances in Econometrics and Modelling, Klumer Academic Publishers, Boston.

Polkovnichenko, Valery, 2003, Household Portfolio Diversification: A Case for Rank Dependent Preferences, Working Paper, University of Minnesota.

Polkovnichenko, Valery, 2004, Limited Stock Market Participation and the Equity Premium, *Finance Research Letters* 1, 24-34.

Poterba, James, and Andrew A. Samwick, 2001, Portfolio Allocations Over the Life Cycle, in Seiritsu Ogura, Toshiaki Tachibanaki, and A. David Wise, eds.: Aging Issues in the United States and Japan, The University of Chicago Press, Chicago, Ill., 65-103.

Quan, Daniel C., and Sheridan Titman, 1997, Commercial Real Estate Prices and Stock Market Returns: An International Analysis, *Financial Analysts Journal* 53, 21-34.

Samuelson, Paul A., 1969, Lifetime Portfolio Selection By Dynamic Stochastic Programming, *The Review of Economics and Statistics* 51, 239-246.

Samuelson, Paul A., 1991, Long-Run Risk Tolerance When Equity Returns Are Mean Regressing: Pseudoparadoxes and Vindication of "Businessman's Risk," in William C. Brainard, William

D. Nordhaus and Harold W. Watts, eds.: Money, Macroeconomics, and Economic Policy: Essays in Honor of James Tobin, MIT Press, Cambridge, Mass..

Santos, Tano, and Pietro Veronesi, 2004 Labor Income and Predictable Stock Returns, Working Paper, Columbia University and University of Chicago.

Schroder, Mark, and Costis Skiadas, 1999, Optimal Consumption and Portfolio Selection with Stochastic Differential Utility, *Journal of Economic Theory* 89, 68-126.

Schroder, Mark, and Costis Skiadas, 2003, Optimal Lifetime Consumption-Portfolio Strategies under Trading Constraints and Generalized Recursive Preferences, *Stochastic Processes and their Applications* 108, 155-202.

Schroder, Mark, and Costis Skiadas, 2004, Lifetime Consumption-Portfolio Choice under Trading Constraints, Recursive Preferences and Nontradeable Income, *Stochastic Processes and their Applications*, forthcoming.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, 2001, Asset Pricing with Idiosyncrtic Risk and Overlapping Generations, Working Paper, Carnegie Mellon University, Stockholm University, and the Wharton School.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, 2002, Cyclical Dynamics of Idiosyncratic Labor Market Risk, *Journal of Political Economy*, forthcoming.

Telmer, Chris I., 1993, Asset-Pricing Puzzles and Incomplete Markets, *Journal of Finance* 48, 1803-1832.

Viceira, Luis M., 2001, Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income, *Journal of Finance* 56, 433-470.

Vissing-Jørgensen, Annette, 2002, Limited Asset Market Participation and the Elasticity of Intertemporal Substitution, *Journal of Political Economy* 110, 825-853.

Vissing-Jørgensen, Annette, and Orazio P. Attanasio, 2003, Stock-Market Participation, Intertemporal Substitution, and Risk-Aversion, *American Economic Review* 93, 383-391.

Wachter, Jessica A., 2002, Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets, *Journal of Financial and Quantitative Analysis* 37, 63-91.

Xia, Yihong, 2001, Learning about Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation, *Journal of Finance* 56, 205-246.

Yao, Rui, and Harold H. Zhang, 2003, Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraint, *Review of Financial Studies*, forthcoming.

7 Figures

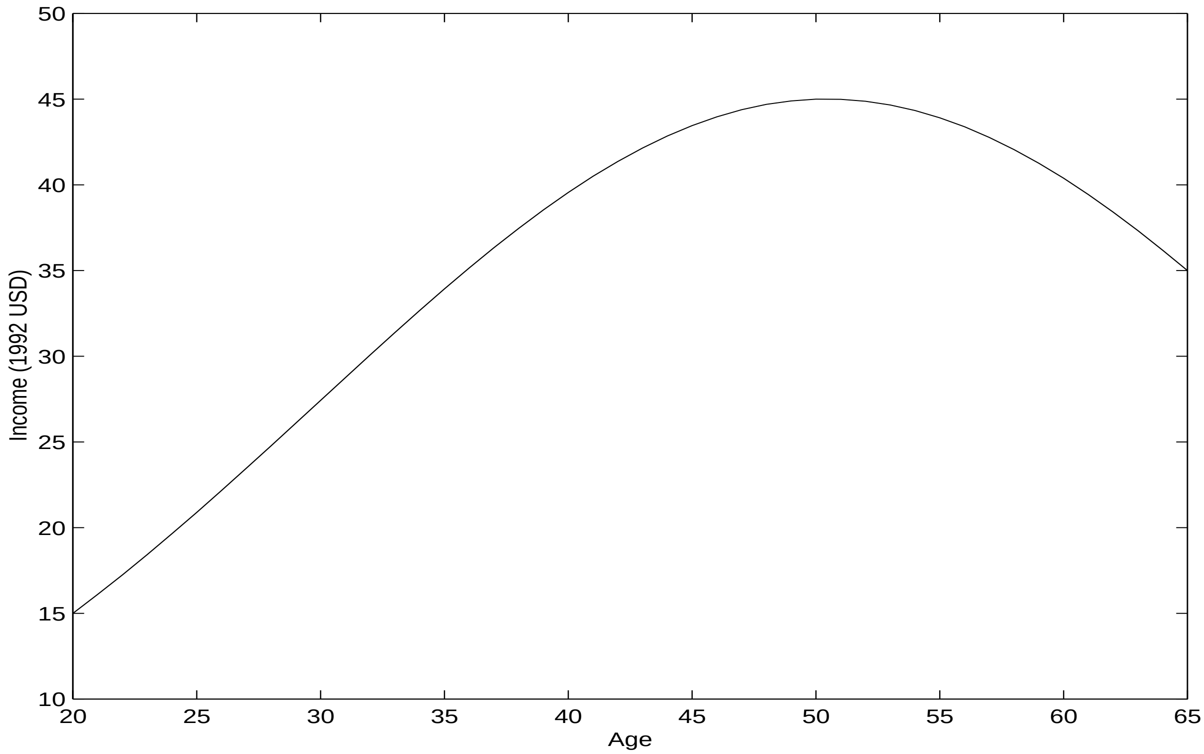

Figure 1: The plot depicts the life-cycle deterministic labor income profile that results from our calibration of the $\alpha(t)$ term in (7). The agent enters the job market at age 20, earning an annual income of \$15,000 in 1992 USD, and retires at age 65.

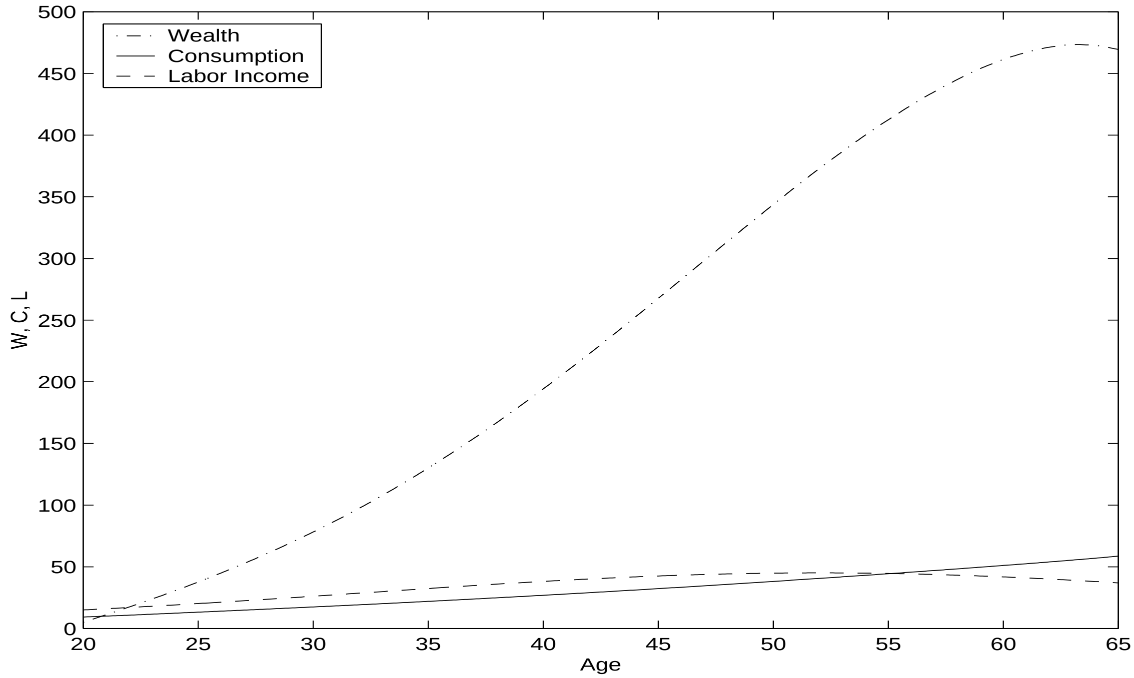

Panel A: Wealth, Consumption, and Labor Income in 1992 USD.

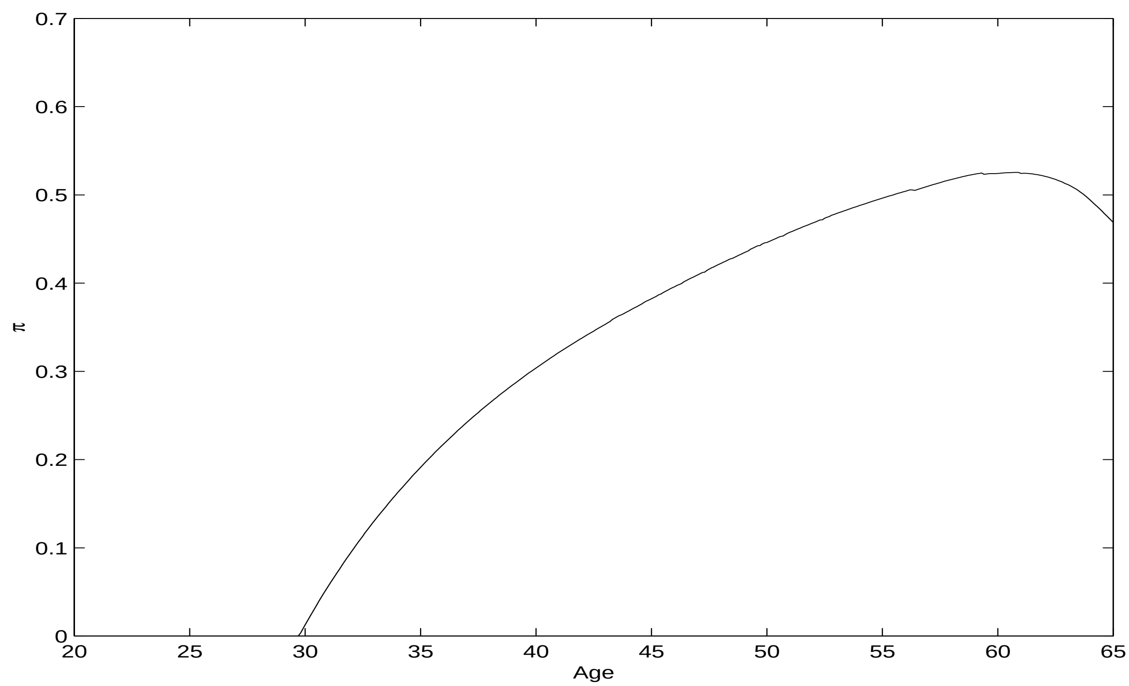

Panel B: Stock Holdings.

Figure 2: Life-cycle profiles of wealth, consumption, labor income, and stock holdings for the baseline case parameters.

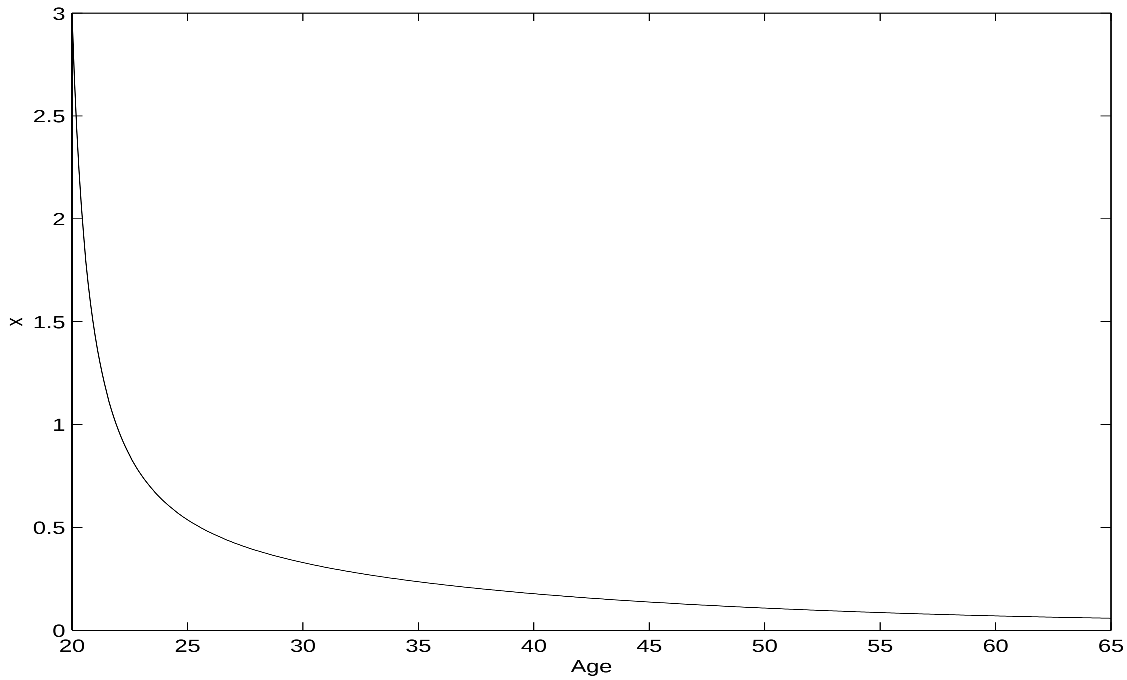

Panel A: Life-cycle profile of the X state variable.

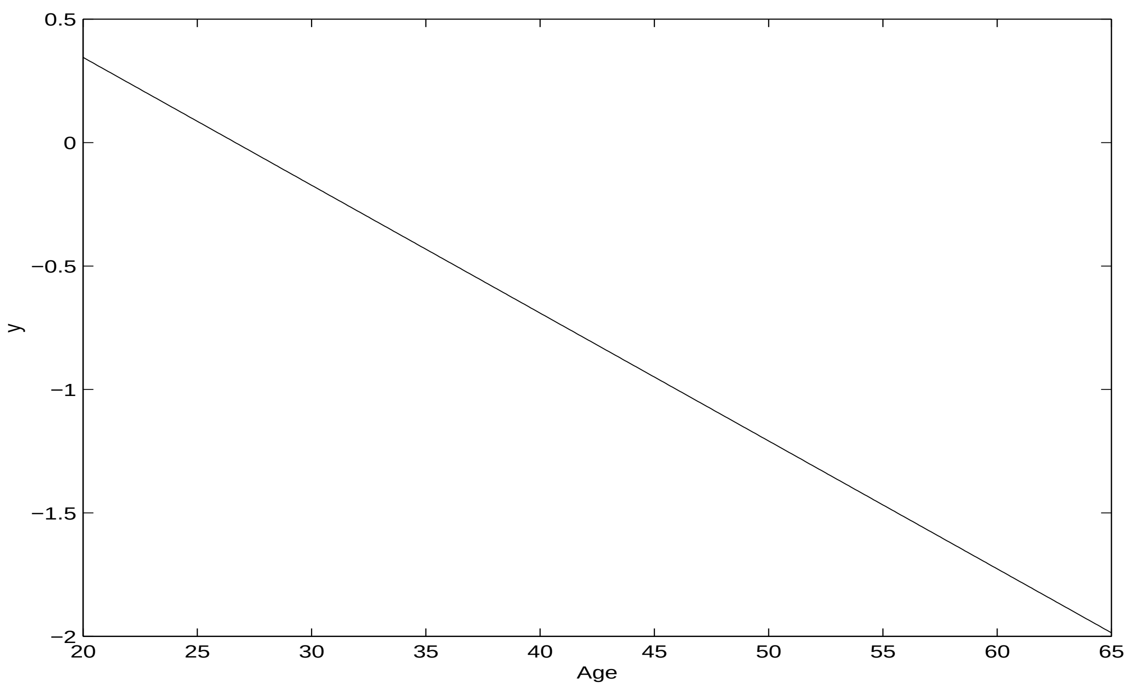

Panel B: Life-cycle profile of the y state variable.

Figure 3: Life-cycle profiles of the model state variables for the baseline case parameters.

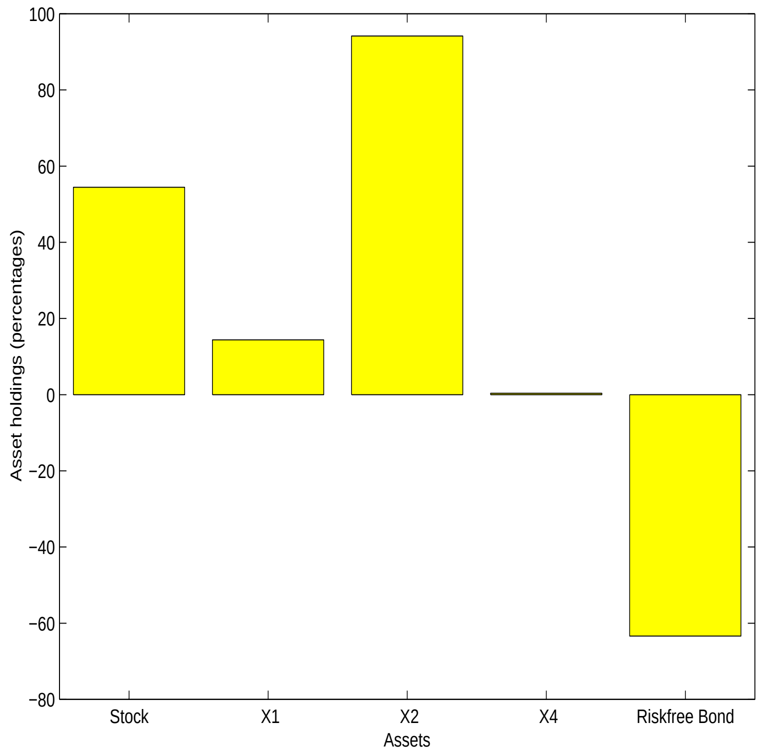

Figure 4: Decomposing the value of human capital into its various components. The graph shows the proportions invested in various securities (stock market S, pseudo-securities X_j, $j = 1, 2$, and 4, and risk-free money market B) that replicates the long position in human capital (i.e., the present value of future labor income flows).

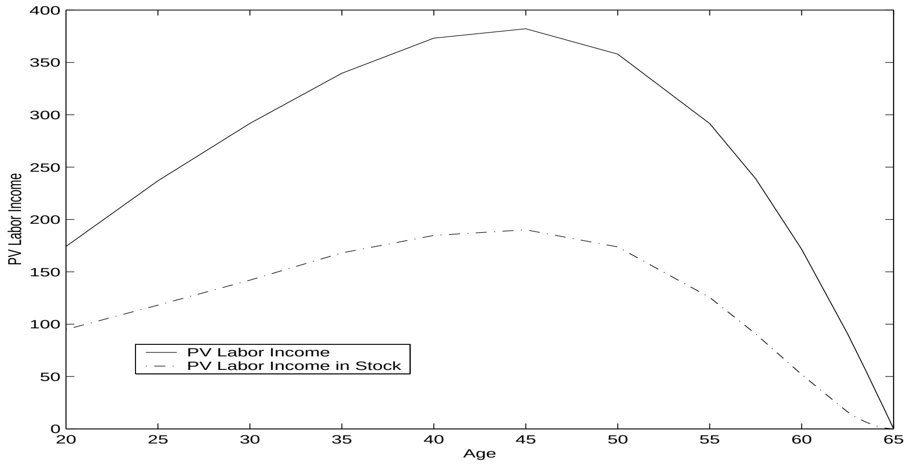

Panel A: Present value of labor income and present value of labor income tied up in stock.

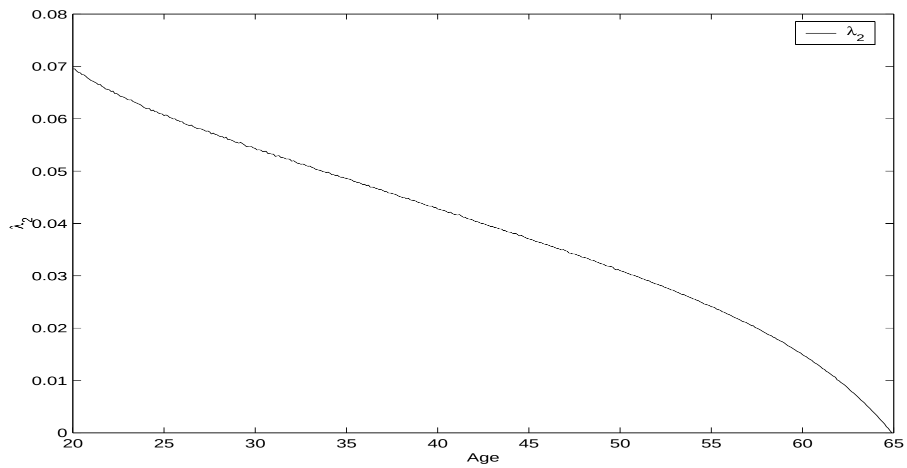

Panel B: Risk premium on the pseudo-security X_2.

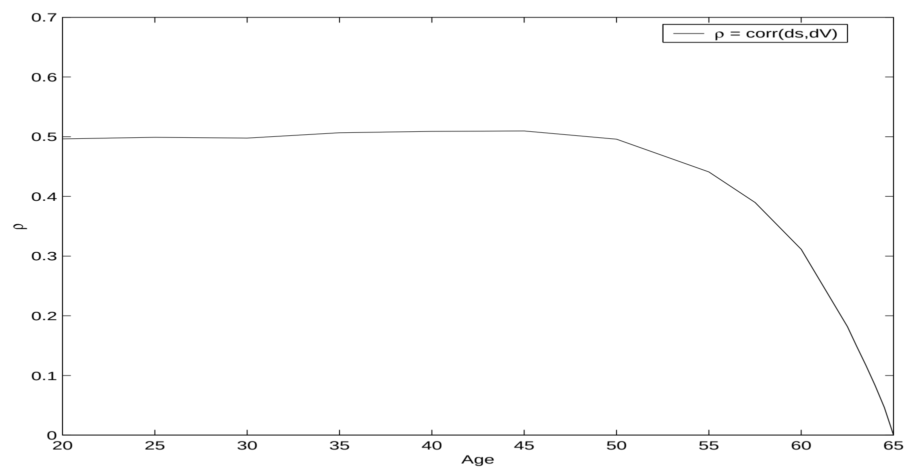

Panel C: Correlation of stock returns and returns to human capital.

Figure 5: The properties of human capital for the baseline case parameters.

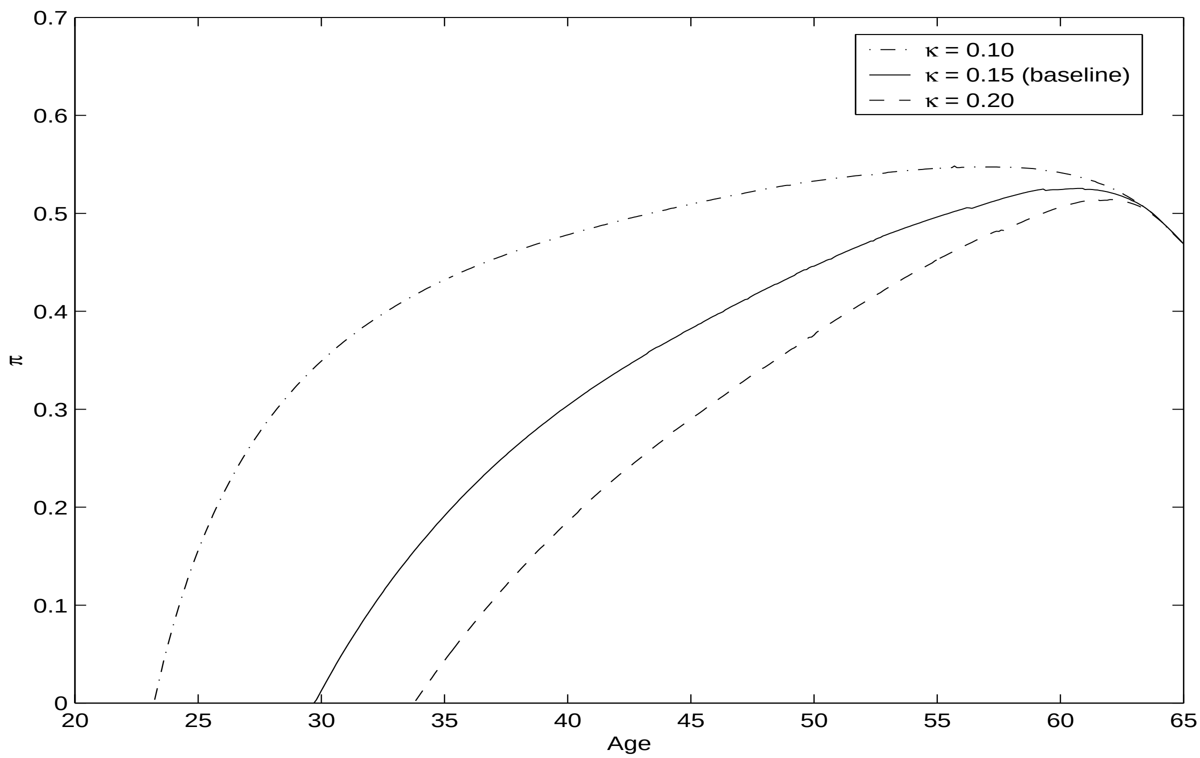

Figure 6: Life-cycle profiles of stock holdings. Sensitivity to the κ coefficient.

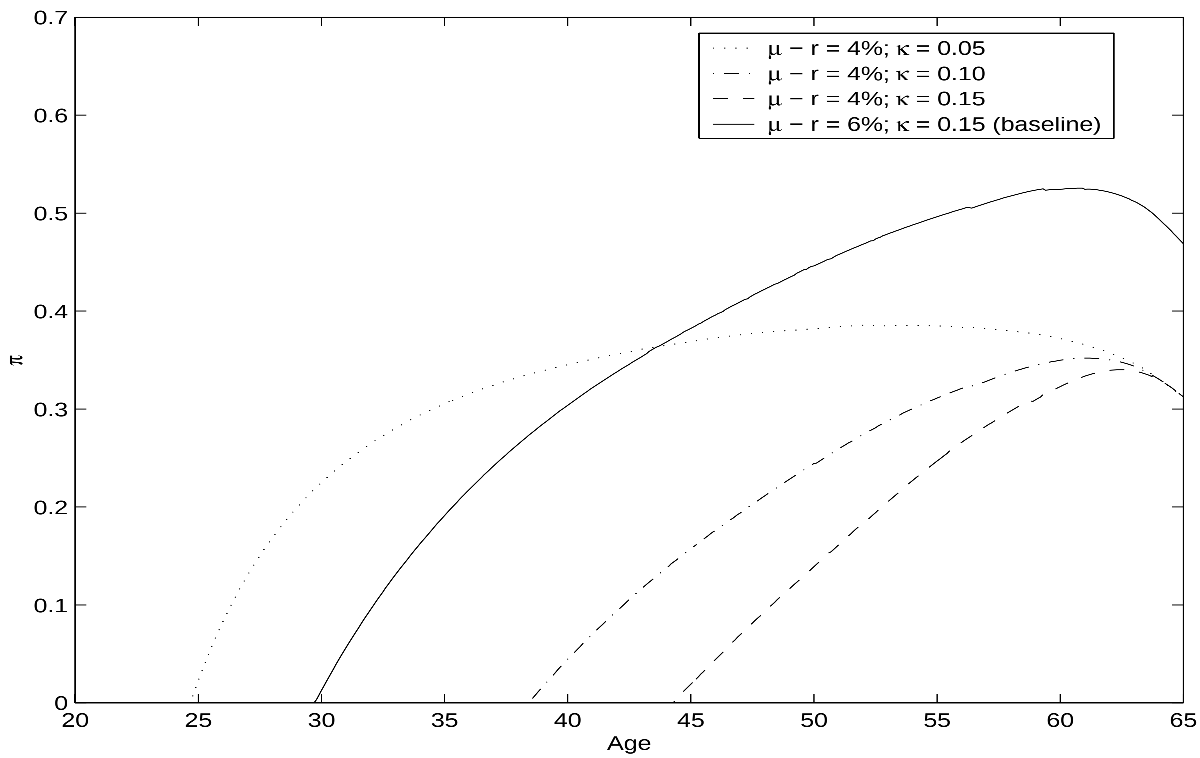

Figure 7: Life-cycle profiles of stock holdings. Sensitivity to the κ coefficient and the risk premium.

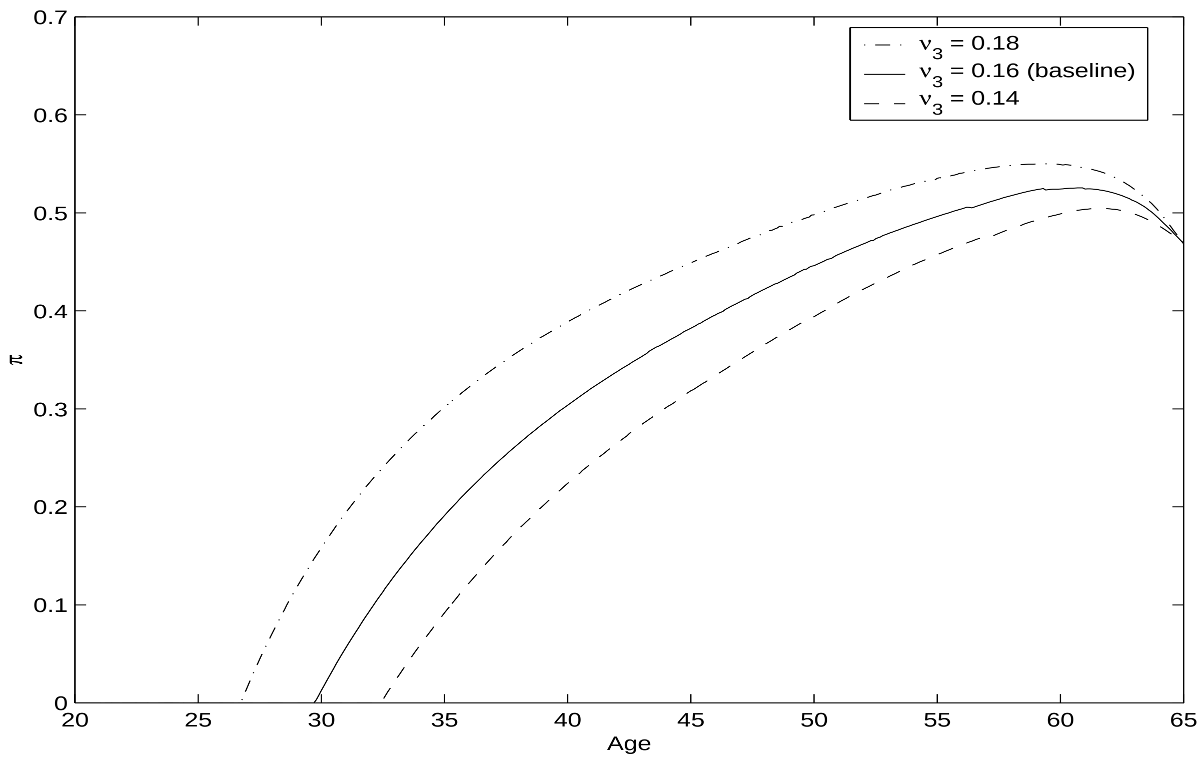

Figure 8: Life-cycle profiles of stock holdings. Sensitivity to the ν_3 coefficient.

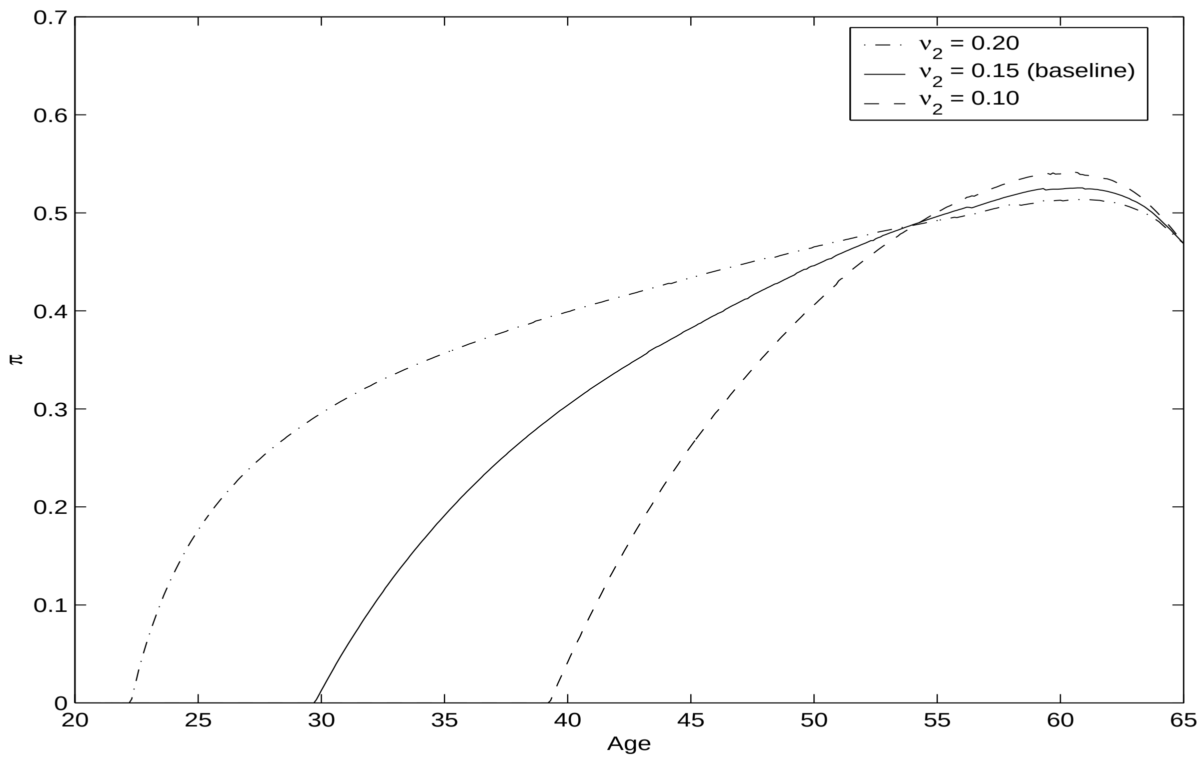

Figure 9: Life-cycle profiles of stock holdings. Sensitivity to the ν_2 coefficient.

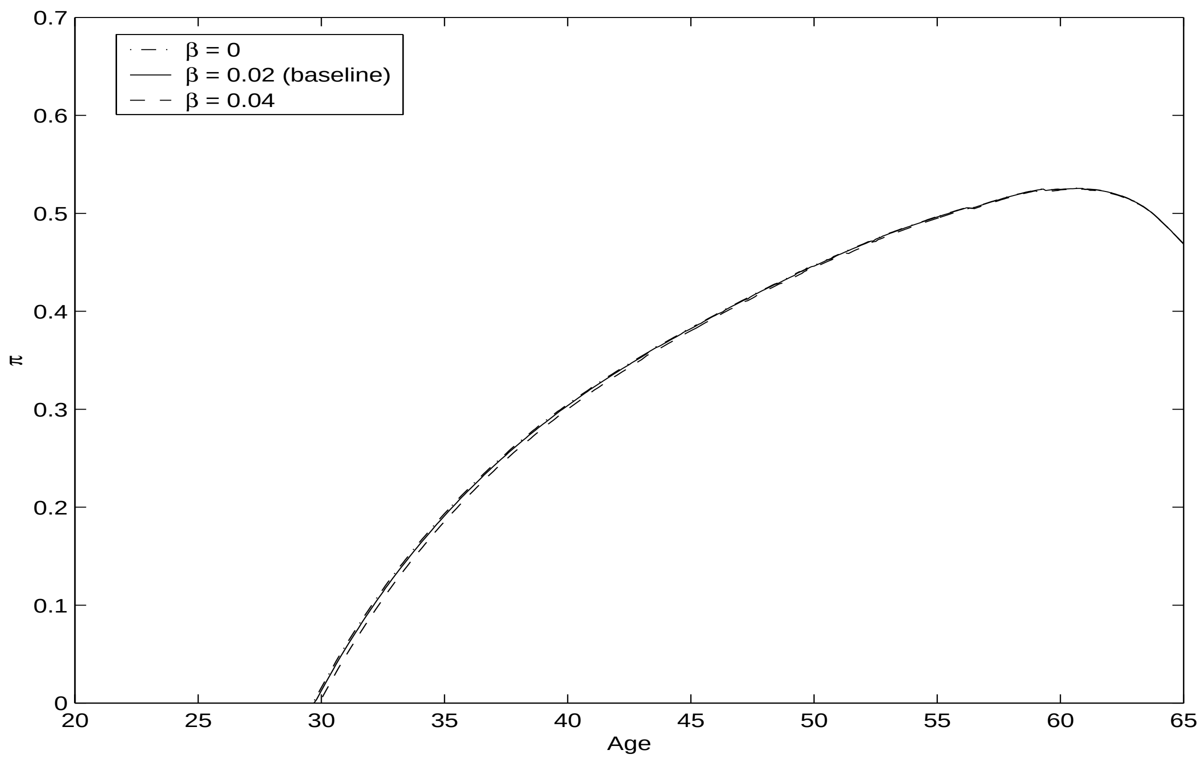

Figure 10: Life-cycle profiles of stock holdings. Sensitivity to the β coefficient.

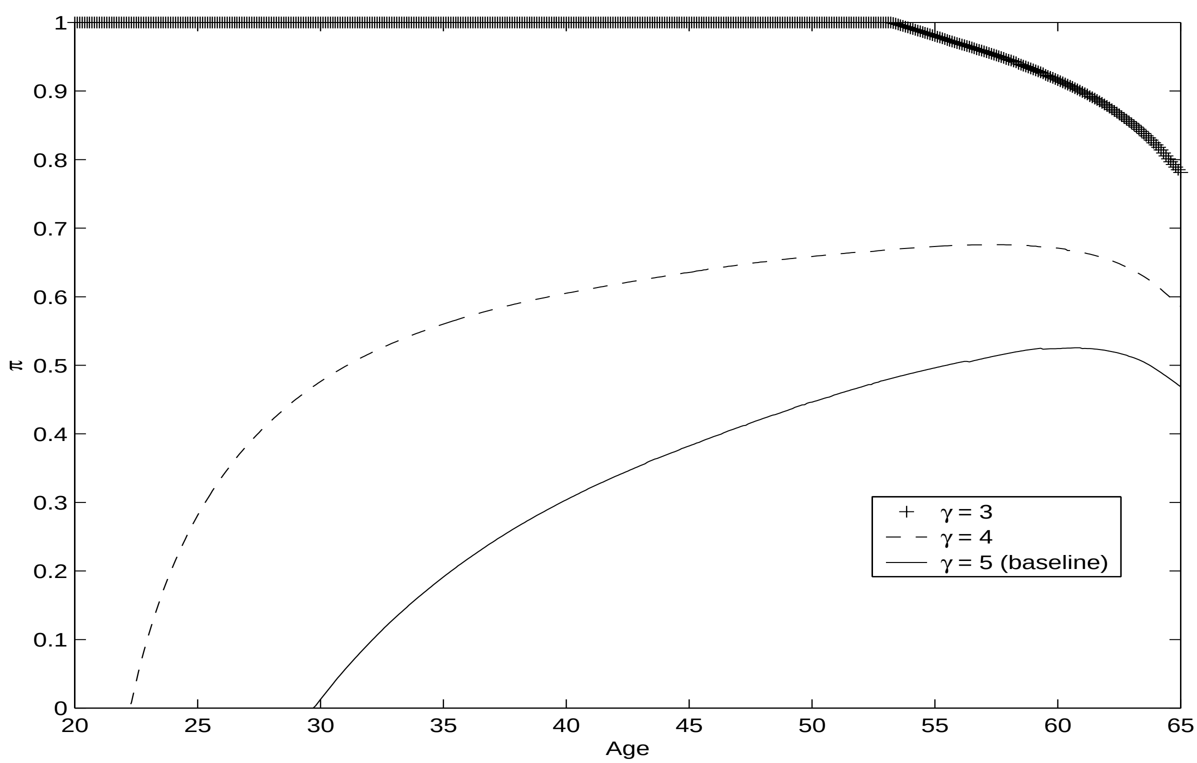

Figure 11: Life-cycle profiles of stock holdings. Sensitivity to the relative risk aversion γ coefficient.

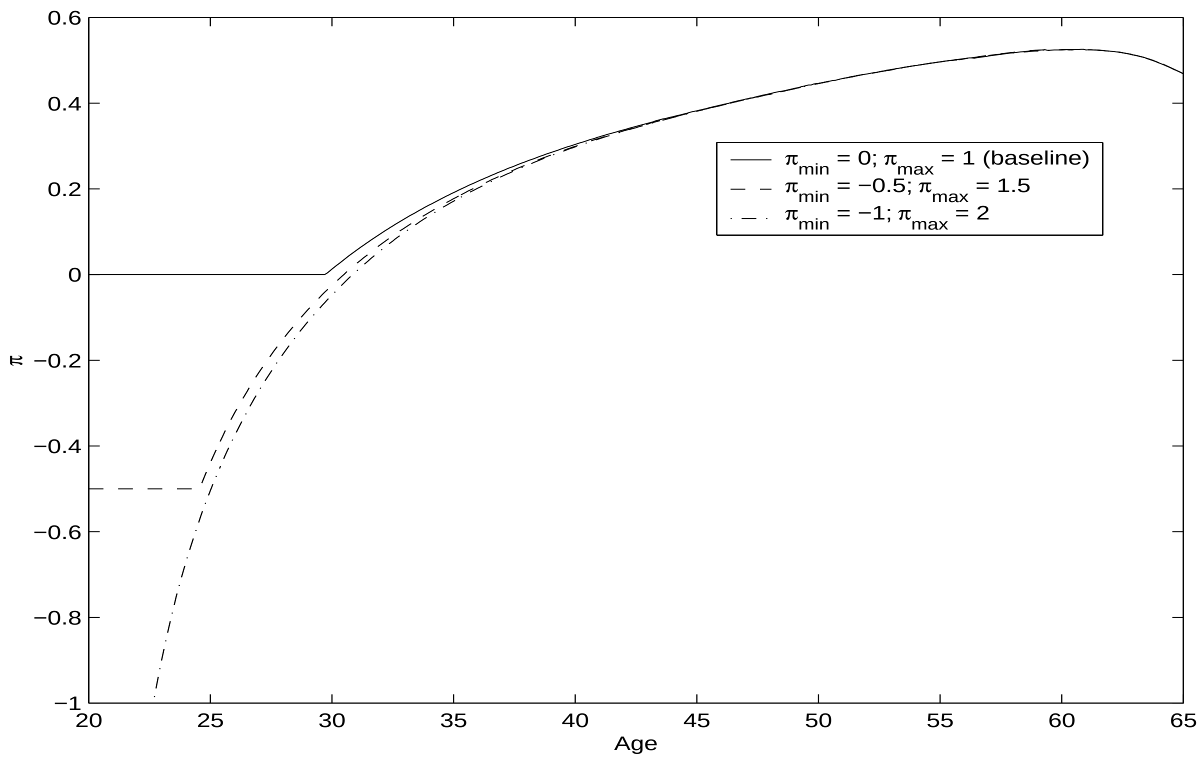

Figure 12: Life-cycle profiles of stock holdings. Sensitivity to the short sale constraints.