Network for Studies on Pensions, Aging and Retirement

NETSPAR ACADEMIC SERIES

Inflation, Money Demand and Portfolio Choice

Kosuke Aoki, Alexander Michaelides and Kalin Nikolov

DP 02/2014-101

Inflation, Money Demand and Portfolio Choice*

Kosuke Aoki†
University of Tokyo

Alexander Michaelides‡
Imperial College Business School, University of Cyprus,
CEPR, CFS and NETSPAR

Kalin Nikolov§
European Central Bank

First Draft: September 2012
Current Draft: February 2014

*We thank the Central Bank of Cyprus, the Bank of England and the European Central Bank for hospitality when this paper was written. This paper has been prepared by the authors under the Wim Duisenberg Research Fellowship Program sponsored by the ECB. Any views expressed are only those of the authors and do not necessarily represent the views of the ECB or the Eurosystem. For helpful comments and suggestions we thank Alexander Bick (discussant), Josef Bonaparte, Joao Cocco, Russell Cooper, Luigi Guiso, Michael Haliassos, Tullio Japelli, Valentina Michelangeli (discussant) and seminar participants at the ECB, the European Conference on Household Finance, EEA, EFMA, Imperial College Business School, Queen Mary, SED, Swedish School of Economics, University of Miami, and University of Warwick. We are responsible for any remaining errors.

†Faculty of Economics, University of Tokyo, 7-3-1 Hongo, Bunkyo, Tokyo, 113-0033, Japan. E-mail: kaoki@e.u-tokyo.ac.jp.

‡Department of Finance, Imperial College Business School, South Kensington Campus, SW7 2AZ. Tel: +44(0)2075949177. E-mail: a.michaelides@imperial.ac.uk.

§Research Department, European Central Bank, Kaiserstrasse 29, Frankfurt, Germany. E-mail: kalin.nikolov@ecb.int.

Abstract

We introduce a money demand motive in a life-cycle portfolio choice model and estimate the structural parameters that can generate limited stock market participation and plausible holdings of money, bonds and stocks. The model predicts an increase in bond holdings over the life cycle, and a declining share of money in portfolios as wealth increases. Both predictions are consistent with the data, even though the model overpredicts (underpredicts) stock (bond) holdings in early life. When mean inflation approaches zero, the share of money in the financial portfolio rises at the expense of both bond and stockholdings, generating simultaneously a lower stock market participation rate.

JEL Classification: E41, G11.

Key Words: Life Cycle Models, Portfolio Choice, Inflation, Money Demand, Stock Market Participation, Uninsurable Labor Income Risk.

1 Introduction

In the recent large literature on portfolio choice,[1] households are assumed to choose between different real assets (typically bonds, stocks and/or housing), ignoring the fact that all transactions in the data are actually done in nominal terms. These models therefore cannot study the effects of inflation on household money demand, portfolio choice and stock market participation decisions. On the other hand, the monetary economics literature starts out with a nominal model and inflation has important real effects and implications for policy. Nevertheless, canonical models in the money demand literature follow the Baumol-Tobin analysis and typically focus on the distinction between money and bonds as a proxy for all other assets in the household portfolio (see, for example, Alvarez and Lippi (2009) and Lippi and Secchi (2009)).

However, a potentially more important decision (especially for the richest part of the population) involves the asset allocation between transaction-type balances (like money) and stocks (that have a substantially different risk-return trade-off from bonds). Moreover, as documented by Doepke and Schneider (2006), inflation changes the real value of nominal assets causing large redistributionary effects in the macroeconomy. Thus, both the level and volatility of inflation have the potential of affecting household consumption and portfolio choice by changing the real returns of different assets.

To build an empirically relevant model of household money demand and portfolio choice, we first examine the empirical properties of life-cycle household financial decisions across three broad asset classes: money (transaction accounts), bonds and stocks. We use both the 2001 and 2007 U.S. Surveys of Consumer Finances (SCF) for this purpose. In the 2001 (2007) SCF 91% (92%) of all households have a transactions/liquidity account, 63% (64%) have

[1] See Campbell (2006) for a recent excellent survey.

a positive amount of bonds and 52% (50%) participate in the stock market (including participation through retirement plans).

In the SCF data, wealthier households tend to hold a mixture of money, bonds and stocks with money always featuring in the portfolio. Nevertheless, the proportion of money as a percentage of total financial assets is smaller for richer rather than poorer households. More precisely, households that do not participate in the stock market tend to be poorer and tend to hold a very large proportion (more than two thirds throughout their life cycle) of their financial wealth in money-like, transaction accounts.

To analyze the effect of inflation on the real economy and asset allocation choices we first need to introduce a role for money balances in a model. Introducing money can vary in complexity from the decentralized search Kiyotaki and Wright (1989) setup, to other money demand models, such as cash-in-advance (Lucas and Stokey (1987)), money-in-the-utility function (Sidrauski (1967)) and shopping-time approaches (McCallum and Goodfriend (1987)).

Given that our purpose is to develop a tractable model that can be confronted with the data, we use an approach similar to shopping time models.[2] Specifically, we assume that money provides liquidity services and therefore a higher amount of money lowers the cost from having to undertake a given transaction for consumption purposes, other things being equal. Everything else we assume is similar to recent life-cycle models that feature intermediate consumption and stochastic uninsurable labor income in the tradition of Deaton (1991) and Carroll (1997)[3], and as extended in the life cycle portfolio choice literature by Cocco, Gomes, and Maenhout (2005), for instance. One

[2] For recent applications of shopping time models, see, for example, Mulligan and Sala-i-Martin (2000).

[3] Attanasio, Banks, Meghir, and Weber (1999), Gourinchas and Parker (2002) and Cagetti (2003) extend this tradition and estimate the structural parameters of life cycle models with a single real asset (a riskless bond).

nice feature of our setting is that it nests the life cycle portfolio models where bonds and stocks are real assets and money does not circulate in the economy. We also introduce a fixed cost to participate in the stock market[4] to generate low stock market participation for one group of households.

Ex ante heterogeneity in risk aversion and the elasticity of intertemporal substitution is used to generate poorer and richer households, as in Gomes and Michaelides (2005). Given the lack of guidance in picking the shopping technology parameters, we estimate these using a method of simulated moments (MSM) estimation technique. We estimate the structural parameters by matching moments from the 2007 Survey of Consumer Finances data. Specifically, we simulate the model to match mean financial wealth to labor income over the life cycle for stockholders and non-stockholders, and the portfolio shares across money, bonds and stocks for stockholders and between money and bonds for non-stockholders. We also estimate the one-time fixed cost to generate the 50% participation rate in the 2007 data.

The estimated preference parameters are reasonable. The annual discount factor is estimated at 0.93 and there is a relatively strong bequest motive for stockholders (consistent with De Nardi, French, and Jones (2010) who find a stronger bequest motive for wealthier households that are more likely to participate in the stock market). The estimated shopping cost parameters for participants and non-participants generate low shopping costs over the life cycle, ranging between half and two percent of mean labor income. The fixed cost that generates stock market non-participation is estimated at 3.8% of mean labor income, an estimate that is in line with previous calibration

[4]This follows a large recent literature on household portfolio choices to generate limited stock market participation, see for example, Vissing-Jorgensen (2002), Haliassos and Michaelides (2003), Gomes and Michaelides (2005), Alan (2006), Paiella (2007), Attanasio and Paiella (2010), Bonaparte, Cooper, and Zhu (2012), Alvarez, Guiso, and Lippi (2012) and Cooper and Zhu (2013).

work (for instance, Gomes and Michaelides (2005)) and estimation work (for example, Bonaparte, Cooper, and Zhu (2012)).

The estimated model matches quantitatively limited stock market participation, and the share of wealth in money, bonds and stocks over the later parts of the life cycle. On the other hand, the model overpredicts (underpredicts) the share of wealth allocated to stocks (bonds) in the early part of the life cycle, even though it matches well the share of wealth allocated to liquid balances (money) in that part of the life cycle.

Moreover, the model can replicate the fact that the share of wealth in stocks increases as financial wealth rises over some parts of the life cycle. This prediction is consistent with the data, while this stylized fact has been a priori inconsistent with recent models of household portfolio choice, as pointed out by Wachter and Yogo (2010). The young (who start out with low financial wealth) have higher liquidity needs and economize on shopping costs by holding liquid balances in the form of money. As financial wealth increases over the life cycle, they diversify into stocks and then hold a mixture of all three assets later on in the life cycle.

We next use the estimated model to provide answers to interesting counterfactual questions. What are the effects of both the level and volatility of inflation on money demand and asset allocation? For stockholders (nonstockholders) higher mean inflation causes a reallocation away from money into stocks (bonds). The reallocation from money into stocks is stronger for younger stockholders who have a stronger demand for stocks in the model and are more likely to increase it when the risk-return trade-off is improved. In addition, younger households devote a higher share of their financial wealth in liquid balances and therefore are more keen to reallocate out of money in the presence of higher inflation.

Perhaps equivalently, setting inflation to zero increases dramatically the demand for money. Bonds decline as a share of household portfolios. This happens for both stockholders and non-stockholders but is quantitatively stronger for the latter. Interestingly, stock market participation is also substantially reduced under zero inflation because money improves as a store of value.

What are the hedging demands generated by inflation? We find that hedging demands are small for low to moderate rates (10%) of inflation when inflation is perceived to be i.i.d. (as in recent years). We find this surprising but, given the level of idiosyncratic uncertainty faced by households, we think it is a reasonable conclusion. Further research can ascertain whether the introduction of a persistent inflation process or money illusion can change this conclusion.

In terms of the literature, we view the paper as contributing towards understanding money demand and portfolio choice in the presence of nominal assets. Typically, research on money demand focusses on the distinction between money and safe bonds (see, for example, Mulligan and Sala-i-Martin (2000), Alvarez and Lippi (2009)). In our model we make explicit the choice between money (that earns a zero nominal return) and other assets like bonds and stocks that earn the historically observed rates of return. Moreover, we estimate the structural parameters of a life cycle model that can replicate the observed portfolio choices.

The other strand of the literature that the model relates to is the recent life cycle saving and portfolio choice literature (Cocco, Gomes, and Maenhout (2005), Gomes and Michaelides (2005), Polkovnichenko (2007) and Wachter and Yogo (2010) to name some examples). In all these papers, however, the choice is between real assets (real bonds and real stocks) and therefore the effects of inflation on consumption, wealth accumulation, money demand and portfolio choices cannot be analyzed.

The rest of the paper is organized as follows. Section 2 discusses some stylized facts regarding money holdings over the life cycle. Section 3 presents the model, and Section 4 reports the estimated parameters. Section 5 presents the benchmark numerical results and Section 6 conducts several comparative statics. Section 7 concludes.

2 Empirical Evidence on Life Cycle Asset Allocation and Participation in Different Asset Markets

We first need a working definition of money to proceed. We use the same definition that the SCF uses to construct the variable LIQ in the public extract of the data set. Specifically, LIQ is defined as the sum of all checking, saving, money market, deposit and call accounts.

We focus on using a single cross section, the 2007 Survey of Consumer Finances, for establishing the stylized facts about the holdings of money in household portfolios over the life cycle. We have repeated the analysis below for all triennial surveys between 1989 and 2007 and we can report that the results from the 1998, 2001, 2004 and 2007 surveys are very similar along the dimensions we report below. Earlier surveys (the 1989 for example) feature lower stock market participation and higher shares of money in the portfolios. The rise of the equity culture in the 1990s is probably responsible for this change. This points towards having to come up with identifying assumptions to decompose the cross sectional results into age, time and cohort effects in this earlier period, if one is interested in secular changes since 1989. Instead of following this approach, we compare the results across the 1998 and 2007 surveys and find that our stylized facts are robust both qualitatively and quantitatively across these four surveys. Cohort effects seem to be less important in this period and we therefore interpret the cross sectional evidence as the set of life cycle facts a good monetary model will need to explain. We therefore leave to future work the time-age-cohort decomposition that could be quite important in understanding the evolution of money demand and stock market participation in the last three decades.

We also restrict our analysis to workers with positive earnings over the year and exclude people with interest in business equity. We focus solely on workers because entrepreneurs have the option of investing in their own business equity (Heaton and Lucas (2000a)) and our labor income process is based on individual labor income data. In the data, most households have a liquid account to undertake their transactions. In the 2001 (2007) SCF 91% (92%) of all households had a transactions/liquidity account, 63% (64%) had a positive amount of bonds and 52% (50%) participated in the equity market (including participation through retirement plans).

One of the well-known stylized facts in the life cycle portfolio choice literature is that financial wealth is correlated with stock market participation (see Campbell (2006) for a recent survey). We estimate the mean amount of financial wealth for households who hold no stocks but just liquid accounts and bonds (in 2007 terms). The mean (median) amount of financial wealth for this group equals 23,512 (2,020) US$, whereas for the stockholders mean (median) financial wealth equals 290,206 (85,300) US$ illustrating the stark dichotomy between households that hold stocks and households that do not. Table 1 reports the levels of financial assets across the two groups over five broad age categories (four during working life and one during retirement) and

Life Cycle Financial Wealth Accumulation

Age Group	Mean (Median) Non-Stockholders	Mean (Median) Stockholders
20-34	6034 (1100)	56966 (20750)
35-45	10688 (1200)	158989 (62550)
46-55	21334 (1800)	275634 (106580)
56-65	31433 (3540)	434306 (155600)
66+	47852 (6900)	516218 (177200)

Table 1: Mean (median) financial wealth for the two main groups (non-stockholders/stockholders from the 2007 SCF data. The precise definitions for the different variables are in Appendix A.

shows that stockholders are richer.

Recent work explains this fact using a fixed cost to prevent households from participating in the stock market[5]. We also follow the fixed cost approach to generate stock market non-participation in the structural model.

A second well known issue in the literature that comes out from Table 1 is the skewed distribution of financial wealth which affects the choices researchers need to make when bringing models to the data. In general, there are three main mechanisms being used to match the observed wealth distribution: heterogeneous discount rates (Krusell and Smith (1998)), bequests (De Nardi (2004)), and a combination of bequests and entrepreneurship (Castaneda, Diaz-Gimenez, and Rios-Rull (2003)). These are general equilibrium models with a single asset, whereas we want to eventually solve a model with three different assets and different rates of return. Rather than complicating the model further we abstract from matching the wealth distribution exactly. Instead we focus on matching the evolution of mean financial wealth to labor income over the life cycle. Eventually a general equilibrium model can be calibrated to match these magnitudes, after the demands for different assets have

[5] This follows a large recent literature on household portfolio choices to generate limited stock market participation, see for example, Vissing-Jorgensen (2002), Haliassos and Michaelides (2003),Gomes and Michaelides (2005), Alan (2006), Paiella (2007), Attanasio and Paiella (2010) and Bonaparte, Cooper, and Zhu (2012).

Life Cycle Financial Wealth Accumulation Relative to Mean Labor Income

Age Group	Mean (Wealth/Income) Non-Stockholders	Mean (Wealth/Income) Stockholders
20-34	0.31	0.85
35-45	0.26	1.53
46-55	0.56	2.65
56-65	0.77	7.02
66-75	2.48	15.7

Table 2: Mean financial wealth relative to labor income for the non-stockholders/stockholders in the 2007 SCF data. The definitions for the different variables are in Appendix A

been pinned down.[6] We leave again the more ambitious task of matching the wealth distribution in the context of this monetary model to future work.

Table 2 reports the moments we will be matching. It contains similar information to that in Table 1 but mean financial wealth is normalized by labor income. There is a jump after retirement as retirement income drops below mean working life labor income.

It should also be noted that we do not include home equity in our measures of household financial wealth: we want to think of the available financial assets and how they are allocated across liquid balances. In the structural model we will use a data-driven housing expenditures life-cycle function to exogenously subtract the expenditures that go in housing from labour income, as is done in Gomes and Michaelides (2005) and Love (2010).

We next go deeper into the role of money in the household portfolio and how money allocations change over the life cycle. We compute the mean asset allocations across money, bonds and stocks for the households that hold all three assets (stockholders) and also for the households that hold only money and

[6] This is the approach advocated by Heaton and Lucas (2000b). Solving and understanding the intuition behind general equilibrium heterogeneous agent models is difficult. One useful (or intermediate) step involves matching the demand side holding asset returns exogenous and at their observed historical values before a general equilibrium model that matches the data is constructed.

Life Cycle Portfolio Choice

	Non-Stockholders			Stockholders		
Age Group	α_m	α_b	α_s	α_m	α_b	α_s
20-34	75.8	24.2	0.0	28.6	29.8	41.5
35-45	67.9	32.1	0.0	16.7	35.3	48.0
46-55	62.5	37.5	0.0	14.4	37.6	48.0
56-65	59.2	40.8	0.0	13.5	38.7	47.8
66-75	63.3	36.7	0.0	16.2	42.0	41.8

Table 3: Mean financial portfolios for the non-stockholders/stockholders in the 2007 SCF data. The definitions for the different variables are in Appendix A

bonds (who we define as non-stockholders). For non-stockholders, the mean share of financial wealth allocated in liquid balances (money) is 66.6% (with a standard deviation equal to 40.8%), and therefore the remainder (33.4%) is allocated in bonds (with a standard deviation equal to 40.8%). For stockholders, we find that the mean share of wealth in stocks is 45.5% (with a standard deviation of 28.6%), the mean share of wealth in money is 17.6% (with a standard deviation of 20.5%) and finally the mean share of wealth in bonds is 36.8% (with a standard deviation of 28.1%).

Table 3 reports the mean portfolio shares for money (α_m), bonds (α_b) and stocks (α_s) for the five age groups and across stockholders and non-stockholders. By definition, non-stockholders hold no equities and we can observe that their portfolios are heavily dominated by the money accounts, with a slightly declining profile over the life-cycle. The life-cycle profiles for stockholders do not show any substantial variations, even though there is a small tendency for the share of wealth in money balances to decrease over the working life cycle and increase after retirement. What is immediately apparent, however, is that money is a key feature of household portfolios, despite the rate of return dominance of other assets, with all age groups devoting a substantial percentage of their financial wealth to money holdings.

3 The Model

The model is a nominal version of life-cycle models that are extensively used in the household portfolio literature. Agents work while they are young, and receive a pension after retirement. They are subject to uninsurable labor income risk and borrowing constraints. There are three assets in the economy, money, bonds and stocks, and they are traded in nominal terms. In order to introduce money, we extend the model by introducing nominal assets and transaction frictions.

3.1 Preferences

Time is discrete and t denotes adult age which, following the typical convention in the literature, corresponds to effective age minus 19. Each period corresponds to one year and agents live for a maximum of 81 (T) periods (age 100). The probability that a consumer/investor is alive at time $(t+1)$ conditional on being alive at time t is denoted by ξ_t ($\xi_0 = 1$). Finally, the consumer/investor has a bequest motive.

Households have Epstein-Zin-Weil utility functions (Epstein and Zin (1989), Weil (1990)) defined over one single non-durable consumption good. Let C_{it} and X_{it} denote respectively real consumption and nominal wealth (cash on hand) of agent i at time t. Then the real cash on hand is defined as X_{it}/P_t where P_t denotes the price level at time t. The preferences of household i are defined over real consumption by

$$V_{it}^{j}\left(\frac{X_{it}}{P_t}\right) = \left\{ (1-\beta) C_{it}^{1-1/\psi_j} + \beta \left(E_t \left[\begin{array}{c} \xi_t V_{it+1} \left(\frac{X_{it+1}}{P_{t+1}}\right)^{1-\rho_j} \\ +(1-\xi_t)\varphi_j \left(\frac{X_{it+1}}{P_{t+1}}\right)^{1-\rho_j} \end{array} \right] \right)^{\frac{1-1/\psi_j}{1-\rho_j}} \right\}^{\frac{1}{1-1/\psi_j}} \quad (1)$$

where ρ is the coefficient of relative risk aversion, β is the discount factor, and φ_j determines the strength of the bequest motive. Following Vissing-Jorgensen (2002) we assume that different households are heterogeneous in their inter-temporal elasticity of substitution, ψ. Our economy is populated with two equally-sized groups j, respectively, with high (ψ_H) and low (ψ_L) intertemporal elasticity of substitution. We also utilize ex ante heterogeneity in risk aversion (high (ρ_H) and low (ρ_L)) following Gomes and Michaelides (2005) to generate a stronger variation in precautionary saving motives over the life cycle (and therefore more variation in wealth accumulation profiles). Finally, households are heterogeneous in their bequest motive ($\varphi_H > \varphi_L$).

3.2 Labor Income Process

Following the standard specification in the literature, the labor income process before retirement is given by

$$Y_{it} = Y_{it}^p U_{it} \tag{2}$$

$$Y_{it}^p = \exp(f(t, Z_{it}))Y_{it-1}^p N_{it} \tag{3}$$

where $f(t, Z_{it})$ is a deterministic function of age and household characteristics Z_{it}, Y_{it}^p is a permanent component with innovation N_{it}, and U_{it} a transitory component. We assume that $\ln U_{it}$ and $\ln N_{it}$ are independent and identically distributed with mean $\{-.5 * \sigma_u^2, -.5 * \sigma_n^2\}$, and variances σ_u^2 and σ_n^2, respectively. The log of Y_{it}^p evolves as a random walk with a deterministic drift, $f(t, Z_{it})$. For simplicity, retirement is assumed to be exogenous and deterministic, with all households retiring in time period K, corresponding to age 65 ($K = 46$). Earnings in retirement ($t > K$) are given by $Y_{it} = \lambda Y_{iK}^p$, where λ is the replacement ratio ($\lambda = 0.68$).

Durable goods, and in particular housing, can provide an incentive for higher spending early in life. We exogenously subtract a fraction of labor income every year allocated to durables (housing). This empirical process is taken from Gomes and Michaelides (2005) and is based on Panel Study Income Dynamics (PSID) data.

3.3 Specification of shopping cost technology

In order to motivate money holdings, we assume transaction frictions. Our approach is related to shopping time models, first proposed by McCallum and Goodfriend (1987). We modify slightly that specification to incorporate it more easily in the portfolio choice literature. In shopping time models, transaction costs are modeled in terms of foregone time: money can help reduce transaction time. As shown in Lucas (2000), there is a connection between the shopping time models and the inventory-theoretic studies of money (Baumol (1952), Tobin (1956)). More broadly speaking, the transaction cost can include not only a shopping cost but also a cost of selling illiquid assets to finance consumption. Different versions assume different trade-offs in the presence of transactions frictions. For example, Lucas (2000) assumes that agents face a trade-off between hours spent on production and transactions. Ljungqvist and Sargent (2004) (Ch. 24) assume a trade-off between transaction time and leisure.

To generate money holdings, we assume a shopping cost transaction friction (proportional to Y_{it}^{p}), a direct physical cost in consumption goods:

$$\Omega_{it} Y_{it}^{p} = \Omega(C_{it}, M_{it}/P_t; \varepsilon) Y_{it}^{p}, \quad \Omega_C > 0, \quad \Omega_M < 0$$

The cost is increasing in real consumption and decreasing in real money balances. In the benchmark case we assume

$$\Omega_{it} Y_{it}^p = \varepsilon \left(\frac{C_{it}}{M_{it}/P_t} \right) Y_{it}^p, \quad \varepsilon \geqslant 0 \tag{4}$$

It will be convenient later on to divide variables by the permanent component of labor income. In this case, the shopping cost per unit of Y_{it}^p is given by

$$\Omega_{it} = \varepsilon \left(\frac{c_{it}}{m_{it}} \right), \quad \varepsilon \geqslant 0 \tag{5}$$

where $c_{it} = \frac{C_{it}}{Y_{it}^p}$ and $m_{it} = \frac{M_{it}}{Y_{it}^p P_t}$. Our preferred interpretation is that the transaction cost represents an opportunity cost of time which is why we assume that it is proportional to the permanent component of labor income. The functional form (4) is consistent with Lucas (2000) who shows that the implied money demand function is consistent with the demand function of Baumol (1952) and Tobin (1956). The parameter ε measures the severity of transaction frictions. A large ε means it takes more resources to do transactions and it can be different over the life cycle or across agents.

We model transaction costs as a direct physical cost in terms of consumption goods. An advantage of our approach is that we can treat money in exactly the same way as we treat bonds and stocks because there is no additional margin between money holding decisions and leisure (or labor supply) decisions. Therefore our model maintains the basic structure of the models used in the portfolio choice literature, making the model computationally tractable and making the results easily comparable to those obtained in the literature. The presence of the permanent component of income in our shopping cost formulation (4) is consistent with the spirit of the shopping time technology specification which relates the cost of illiquidity to the wage rate faced by the household. Also, our modeling approach maintains the basic properties of the shopping time models — money demand will be increasing in consumption and decreasing in nominal interest rates.

3.4 Financial Assets and Constraints

The agent has options to hold three kinds of assets: fiat money (M_{it}) , nominal bonds (B_{it}) and nominal stocks (S_{it}). We let X_{it} be nominal "cash on hand" that the agent can use for consumption and portfolio decisions. The budget constraint is given by

$$X_{it} = P_t C_{it} + S_{it} + B_{it} + M_{it} + 1_t(.) P_t F Y_{it}^p. \tag{6}$$

where the indicator function $1_t(.)$ becomes one when the fixed cost to participate in the stock market is incurred (equal to $P_t F Y_{it}^p$ in nominal terms).

We assume that the shopping cost is deducted at the beginning of the next period. Then, the evolution of X_{it} is given by

$$X_{it+1} = R_{t+1}^s S_{it} + R_{t+1}^b B_{it} + M_{it} + P_{t+1} Y_{it+1} - P_{t+1} \Omega_{it} Y_{it}^p \tag{7}$$

where R_{t+1}^s and R_{t+1}^b respectively denote the nominal returns of stocks and bonds. Note that the nominal return of fiat money is unity. Finally, Y_{it+1} is real income at time $t+1$.

Following the portfolio choice literature, we prevent households from borrowing against their future labor income (this prevents households from counterfactually leveraging up to invest in the stock market to take advantage of the equity premium). More specifically we impose the following restrictions:

$$B_{it} \geq 0$$

$$S_{it} \geq 0$$

$$M_{it} \geq 0$$

We have one continuous state variable: X_{it} and the control variables are C_{it}, M_{it}, S_{it} and B_{it}.

Moreover, we use the initial wealth distribution from the data to initialize the simulation when households first enter the labor force at age 20. We refer to this initial financial wealth as bequests later on in the paper.

3.5 Normalizing by Prices and Growth

Let lower case letters denote real variables normalized by the permanent component of labor income (Y^p_{it}). For example the normalized real cash on hand is defined as $x_{it} = X_{it}/(Y^p_{it}P_t)$[7]. The evolution of the state variable is then given by

$$x_{it+1} = \frac{r^s_{t+1}}{g_{it+1}} s_{it} + \frac{r^b_{t+1}}{g_{it+1}} b_{it} + \frac{r^m_{t+1}}{g_{it+1}} m_{it} + U_{it+1} - \frac{\Omega_{it}}{g_{it+1}} \tag{8}$$

where

$$r^s_{t+1} \equiv R^s_{t+1}\pi^{-1}_{t+1}, \quad r^b_{t+1} \equiv R^b_{t+1}\pi^{-1}_{t+1}, \quad r^m_{t+1} \equiv \pi^{-1}_{t+1}$$

are respectively the real returns of stocks, bonds and money, where $\pi_{t+1} \equiv P_{t+1}/P_t$ denotes gross inflation, and $g_{it+1} \equiv Y^p_{it+1}/Y^p_{it}$ is the gross growth rate of the permanent component of labor income.

The representation of consumer preferences in terms of stationary (normalized) units is given in Appendix B.

[7] The only lower-case variable not normalized by P_t is consumption (C_{it}) which was in real terms from the beginning.

4 Parameter Estimation

We estimate the structural model using the method of simulated moments (MSM). Given the large number of parameters in the model we will calibrate certain parameters and then estimate the fixed cost, shopping cost and bequest preference parameters, since these are the parameters for which there is no clear guidance in the literature on how to pick them. The calibration for labor income uses the estimates in Cocco, Gomes, and Maenhout (2005) so that $\sigma_u = 0.1$, $\sigma_n = 0.1$, and $\lambda = 0.68$. Given the positive correlation between education, financial wealth and the probability to participate in the stock market, we use the hump shape process for households with a college degree from Cocco, Gomes, and Maenhout (2005). The correlation between permanent idiosyncratic labor income shocks and the stock return (ρ_{sn}) is set to 0.15 based on the relatively scant evidence that exists with regards to this choice as discussed in Gomes and Michaelides (2005).

We will use exogenous processes for stock and bond returns, inflation and the aggregate component of labor income. Given that we calibrate the cross sectional model to decisions taken in 2007, we use the period 1995 to 2008 to compute descriptive statistics and correlations between these variables. In section 6.4 we provide comparative statics experiments that take into account observed uncertainty with regards to these moments based on a longer time series sample of historical experience. Table 4 reports the first two moments that describe the distributions for inflation, real bond returns, real stock returns and real aggregate wage growth.

We assume an i.i.d. process for real asset returns, real wage growth and inflation. As shown in Table 4, stock returns have a mean real return equal to 6.4% and a standard deviation equal to 18%. It is relatively standard in the

Means and Standard Deviations		
Variable	Mean	S. D.
Inflation	2.5	1.0
Real Bond Returns	2.4	2.6
Real Stock Returns	6.4	18.0
Real Wage growth	2.7	2.0

Table 4: We report the means and standard deviations (S.D.) of key inputs in the annual frequency decision model. The bond return is the return on the one-year bond. Details about the data can be found in Appendix A.

household finance literature to assume that due to annual management fees investors earn a lower equity premium than the historically observed 6%-8%. The bond return process is similarly calibrated with a mean return equal to 2.4% and a standard deviation equal to 2.6%.

We also need to take a stance on the correlations across these variables. These are shown in Table 5. Based on our data sample, we set the correlations between bond and stock returns, as well as the correlation between inflation and the real wage growth, to zero since the observed magnitudes over the period of interest are not statistically different from zero. The contemporaneous correlation between inflation and real bond returns is estimated in this sample to be -0.5, and there is a slight positive correlation between inflation and stock returns over this period (we use 0.25 for that correlation). There is also a positive correlation between wage growth and both bond and stock returns and we set that parameter equal to its estimated counterpart at 0.4.

To use the method of simulated moments we need to decide which moments to match. The key variables of interest for our purposes are the mean holdings of financial wealth over the life cycle for stock holders and non-stockholders and the mean participation rate. Conditional on the participation status we then have asset allocations between money, bonds and stocks sorted by age, and for non-stockholders the allocations between bonds and money. This gives

Correlations

Variable	Inflation	Bond Returns	Stock Returns	Wage growth
Inflation	1.0			
Real Bond Returns	-0.5	1.0		
Real Stock Returns	0.25	0.0	1.0	
Real Wage growth	0.0	0.4	0.4	1.0

Table 5: We report the correlation matrix of key inputs in the decision model. All variables are real, and the bond return is the return on the one-year bond. This is for the period between 1995 and 2008. Details about the data can be found in Appendix A.

a total of twenty six moment conditions.

Given the large number of preference parameters we follow the empirical evidence in Vissing-Jorgensen (2002) and the calibration in Gomes and Michaelides (2005) and use two different values for the elasticity of intertemporal substitution, higher for the households more likely to participate in the stockmarket and lower for the rest. We also utilize heterogeneity in risk aversion rather than the discount rate to generate low propensities to save, thereby generating poorer households who are therefore less likely to incur the fixed stock market participation cost. In our baseline estimation we choose $\psi_H = 0.5$, $\psi_L = 0.3$, $\rho_H = 5.0$ and $\rho_L = 1.3$.

Our preference parameters for non-stockholders might appear strange given that they have lower risk aversion (on average) than stockholders. The incentive to save, however, is what matters in incurring the fixed cost to participate in the stock market and this is monotonically related to the precautionary savings motive (Haliassos and Michaelides (2003)). Our calibration is consistent with the low risk aversion parameters for non-participants found in Alan (2006), Paiella (2007) and Attanasio and Paiella (2010). We also set the bequest parameter for the households that are likely not to accumulate substantial financial wealth equal to zero, consistent with De Nardi, French, and Jones (2010) that find evidence for bequest motives in the richer part of the

Calibrated Structural Parameters

Parameter	Value
ψ_H	0.5
ψ_L	0.3
ρ_H	5.0
ρ_L	1.3
φ_L	0.0
σ_U	0.1
σ_N	0.1
ρ_{sn}	0.15

Table 6: Calibrated structural parameters. ψ_H (ψ_L) is the elasticity of intertemporal substitution for the more (less) wealthy households, ρ_H (ρ_L) is the relative risk aversion coefficient for the more (less) wealthy households, σ_U is the standard deviation of the transitory labor income shocks, σ_N is the standard deviation of the permanent labor income shocks and ρ_{sn} is the correlation between the permanent labor income shocks and stock returns.

population but not for poorer households. For compactness purposes Table 6 lists all calibrated parameters and the standard deviations of the permanent and transitory labor income shocks as used by Carroll (1997) and Cocco, Gomes, and Maenhout (2005).

We then estimate the remaining structural parameters for which there is much less guidance from the empirical literature. At this stage we also introduce another layer of estimated heterogeneity to be consistent with the empirical regularity that households that remain poorer during the working part of the life cycle also remain poor during retirement. We therefore estimate different parameters for the shopping cost $\{\varepsilon\}$ and the bequest parameter for the wealthier households φ_H, while keeping the fixed cost parameter F constant across the two groups.[8] We also estimate a constant discount factor (β)

[8] We provide estimates of the structural parameters using Method of Simulated Moments Estimator (MSM) of Duffie and Singleton (1993). The structural parameters collected in a vector $\hat{\theta}$ are determined as:

$$\hat{\theta} = Argmin_{\theta} D' S^{-1} D.$$

Let Y_t and $\tilde{Y}_t$ denote the observations at time t of the actual and simulated endogenous variables, respectively. Let T be the sample size of the observed series whereas $T \cdot H$ data points are simulated to compute moments from the structural model. For the latter, let $Y_{[T]}$

across the two groups.

There is some empirical evidence on the size of the bequest parameter and that this varies across the wealth distribution (De Nardi (2004)), while shopping technology costs should reflect the opportunity cost of time as emphasized by Aguiar and Hurst (2013). On the size of the fixed cost in the context of this model, the closest paper is Alan (2006) and our non-stockholders would be similar to the non-participants modelled through a fixed cost in that paper.

5 Results

The estimated parameters for our model are given in table 7. The results are consistent with previous estimates in the literature. There is some evidence for a bequest motive needed because financial wealth is not fully decumulated during retirement, this is consistent with, among others, De Nardi (2004). There are no estimates for the equivalents of the shopping cost parameters (from microeconomic data) against which we can compare our results (this was also one of the reasons for performing structural estimation). The implied shopping cost varies between 0.5 and 2.0 percent of mean annual labor income that we view as a reasonable transaction cost and is consistent with Lucas (2000).

The fixed cost to generate non-participation is consistent with the low costs

and $\tilde{Y}_{[TH]}$ denote the vectors of actual and simulated endogenous variables of length T and TH, respectively. We have:

$$D = \left(\frac{1}{T} \sum_{t=1}^{T} moments(Y_t) - \frac{1}{TH} \sum_{t=1}^{TH} moments(\tilde{Y}_t) \right).$$

where moments() denotes a particular moment. The asymptotically efficient optimal weighting matrix S^{-1} equals the inverse of the variance-covariance matrix of the data. Following Appendix B in De Nardi, French, and Jones (2010), we use a diagonal weighting matrix for S^{-1} with the elements along the diagonals being the variance of each moment from the data.

Estimated Structural Parameters		
Parameter	Estimate	Standard error
β	0.93	0.01
φ_H	4.2	0.02
ε_H	0.012	0.003
ε_L	0.00006	0.00001
F	0.0385	0.01

Table 7: Estimated structural parameters. β is the annual discount factor, φ_H is the bequest parameter for the wealthier households, ε_H (ε_L) is the shopping cost parameter for the more (less) wealthy households, and F is the fixed cost incurred to participate in the stock market. Standard errors are computed through numerical derivatives.

that have either been calibrated (Gomes and Michaelides (2005)) or estimated (Alan (2006), Bonaparte, Cooper, and Zhu (2012)) to generate stock market non-participation for poorer households. This cost should be interpreted as a short cut for anything ranging from inertia, behavioral biases, low trust in the stock market (Guiso, Sapienza, and Zingales (2008)), observation and transaction costs stemming from rational inattention (Alvarez, Guiso, and Lippi (2012)), or repeated costs from having a stock trading account (small annual trading costs can add up over a few years as in Bonaparte, Cooper, and Zhu (2012)).

The bequest parameter heterogeneity is consistent with recent estimates (De Nardi (2004)) that find a stronger bequest motive for richer households. In the context of this model the stronger bequest motive for richer households reflects the slow decumulation of financial wealth during retirement for this segment of the population. The estimated discount factor is within the range of recent estimates (Attanasio, Banks, Meghir, and Weber (1999), Gourinchas and Parker (2002), Cagetti (2003)) from structural estimation of life cycle consumption models. Finally, the shopping technology estimation generates a higher cost parameter for the richer households than for the poorer ones.

Without this heterogeneity poorer households spend a substantial amount of their small financial wealth in shopping costs, a potentially counterfactual implication. Using this heterogeneity allows us to replicate the selected moments across the two groups without large shopping costs arising from holding money.

What are the policy functions and life cycle profiles implied by these parameter estimates? Figure 1 shows the portfolio choice policy functions of the low risk aversion- low EIS group against the continuous state variable (cash on hand, x). We show the policy functions for the young (age 25), middle-aged (age 55) and retirees (age 85). The solid lines show the portfolio choices of households that have already paid the fixed cost to participate in the stock market. The dashed lines show the choices of non-participants. The vertical axis plots portfolio shares as a percentage of financial wealth invested in each asset (between zero and one due to the no borrowing/no short sale constraints). The three age groups mainly hold money when their cash on hand is small. Especially, the young and the middle-aged agents invest almost their entire assets in money when they are poor. This is consistent with the data that shows that the young and poor agents tend to hold more money. In line with the portfolio literature, the bond share is increasing in cash on hand while the stock share is decreasing in cash on hand.

The dashed lines indicate the policy functions for the non-participants. These households need to accumulate assets in money and bonds before incurring the fixed participation cost to invest in the stock market. They therefore accumulate money first (to minimize the shopping cost), and then bonds and then if they accumulate a sufficient amount of wealth to justify incurring the fixed participation cost, they do so. At that point they reduce their bond holdings and invest in the stock market.

A similar picture arises from Figure 2 that plots the portfolio policy functions for the higher risk aversion - higher EIS group. The first difference from the previous figure is that the participation takes place faster even though these households are more risk averse. Because of a stronger precautionary saving motive they accumulate more saving and have a stronger incentive to participate in the stock market. The second difference is that the share of wealth in stocks is more quickly declining as a function of cash on hand again because of the higher risk aversion.

Figure 3 conditions on participation status (some of the low risk aversion households will have had lucky labor income draws and participate in the stockmarket) and shows the simulated paths of consumption, financial wealth and income over the life cycle. We simulate the model economy with 10,000 individuals in each age cohort starting with the initial financial wealth taken from the 2007 SCF and take the mean of each variable. The non-stockholders do not accumulate much financial wealth, and their consumption tracks labor income over the life-cycle. The stockholders accumulate a higher amount of financial wealth and therefore their consumption is decoupled from labor income.

Finally, Figure 4 shows simulated portfolio choices over the life cycle. The young non-stockholders hold mainly money, and the money share decreases as the agents become older, reaching a minimum at retirement. At that point money holdings increase again as the limited financial wealth is decumulated during retirement. To minimize on shopping costs during retirement, the financial portfolio is re-allocated towards money and bonds are rapidly crowded out from the portfolio.

Stock market participants are richer households that invest heavily in stocks. In the beginning of working life most households keep their wealth in money in order to minimize shopping costs. Only very rich households that have received large bequests[9] find it profitable to pay the fixed cost and invest in stocks. This is why the share of stocks is so high (and the share of money so low) early on in stockholders' lives.

As we can see from Figure 5, by around age 25 other high-saving households (who did not receive a large bequest) accumulate enough wealth to begin investing in stocks (this can be seen from the sharp pick up in the fraction of stock market participants). This leads to a decline in the average wealth to income ratio of stockholders as can be seen from the middle panel of Figure 3. As a result of the entry of poorer households into the stockholder population, the average share of money in stockholders' portfolios rises and the average share of stocks and bonds falls.

As Figure 5 shows, from age 30 onwards, participation does not change as sharply and most of the evolution in stockholder portfolios is driven by changes to existing stockholders' portfolios. The share of wealth allocated to stocks increases until age 40 as households get richer. This occurs because wealth grows much faster than consumption over this period, meaning that households need to devote a much smaller fraction of their portfolios to money holdings while still keeping shopping costs down. Consequently richer households invest more heavily in stocks - a prediction that the model without money typically cannot generate unless extended in other directions (Wachter and Yogo (2010), for example).

Figure 5 plots the profile of the stock market participation rate. This is increasing until age 50 before declining gently over the rest of the life cycle, consistent with the wealth accumulation motive. The average stock participation rate implied by the model is 50.0% versus 50.3% in the 2007 SCF data.

How do other predicted moments compare with the actual ones? We first

[9] We start from an initial wealth distribution taken from the data in the beginning of life. We think of this initial wealth distribution as arising from a bequest motive.

Mean Financial Wealth/Income: Data versus Model				
Age Group	Non-Stockholders		Stockholders	
	Data	Model	Data	Model
20-34	0.31	0.11	0.85	1.34
35-45	0.26	0.08	1.53	2.18
46-55	0.56	0.08	2.65	3.88
56-65	0.77	0.04	7.02	5.80
66-75	2.48	0.03	15.7	8.34

Table 8: Actual versus predicted moments for mean financial wealth relative to mean labor income for the nonstockholders/stockholders. The model is compared to the 2007 SCF data. The definitions for the different variables are in Appendix A

go through the mean wealth to mean labor income ratios which are given in Table 8. The model captures the fact that stockholders are considerably richer than non-stockholders. We observe that the model predicts some underaccumulation in wealth levels relative to the data for both groups but especially for non-stockholders. This wealth underaccumulation was not heavily penalized by the MSM criterion function due to the high variance of the moments in the data. This is especially true for older age groups (65+ in particular). Overall, given the cross-sectional uncertainty in the data, we view the model implied wealth profiles as a reasonable match for the data, even though further work is needed to bring simulated and actual data closer together.

We next present the moments for the portfolio shares. We start with the profiles for non-stockholders, given in Table 9. The model matches well the share of wealth in money in the early parts of the life cycle (ages 20-54) relative to the data, but overpredicts relative to the data from that point onwards. In particular, the share of wealth in bonds is rising over the life cycle in both data and model and the quantitative magnitudes are very similar until around age 54.

Next we move on to the comparison between data and model for stock-

Life Cycle Portfolio Choice by Age				
	Non-Stockholders			
Age	Data	Model	Data	Model
Group	α_m	α_m	α_b	α_b
20-34	75.8	77.9	24.2	22.1
35-44	67.9	57.6	32.1	42.4
45-54	62.5	58.3	37.5	41.7
55-64	59.1	78.9	40.9	21.1
65+	63.3	98.4	36.7	1.6

Table 9: Actual versus predicted moments for mean financial portfolios for the nonstockholders. The model is compared to the 2007 SCF data. The definitions for the different variables are in Appendix A

holders (Table 10). Here, the main failure of the model is that it cannot match bond-holdings for the first two age groups. Equivalently, as in the real model without time-varying investment opportunities, the model overpredicts the share of wealth invested in the stock market. Specifically, the share of wealth in stocks is 64.6% versus 41.5% in the data for the 20-34 age group, and for the next age group (35-44) this share rises to 72.3% in the model versus 48.0% in the data. Note, however, that the share of wealth in stocks can be rising until age 44 according to the model because of the decline in the share of money as financial wealth is accumulated. This is an implication that cannot be generated in the special case of this model when money does not circulate in the economy.

The results illustrate the strong demand for stocks early in life as labor income is mostly seen like a riskless asset. Nevertheless, money is held in the portfolio for transaction purposes, especially for the non-participants in the stock market. Moreover, the share of wealth in money does drop over the life cycle for both groups, both in the model and in the data. Note that the presence of money does change the composition of the portfolio relative to other models in the portfolio choice literature that lump money and bonds in the

Life Cycle Portfolio Choice by Age						
	Stockholders					
Age	Data	Model	Data	Model	Data	Model
Group	α_m	α_m	α_b	α_b	α_s	α_s
20-34	28.6	34.4	29.9	1.0	41.5	64.6
35-44	16.7	18.8	35.3	8.9	48.0	72.3
45-54	14.4	12.9	37.6	31.8	48.0	55.3
55-64	13.5	11.8	38.7	43.4	47.8	44.8
65+	16.2	10.6	42.0	46.5	41.8	42.9

Table 10: Actual versus predicted moments for mean financial portfolios for the stockholders. The model is compared to the 2007 SCF data. The definitions for the different variables are in Appendix A

same category. Specifically, in these models the standard prediction is that stockholders should allocate all financial wealth in stockholding while in this setup the very young (ages 25-44) allocate between 64 to 79 percent of their financial wealth in stocks. The allocation to bonds is still underpredicted relative to the data but we view the model as getting one step closer to matching observed behavior in the data.

6 Comparative Statics

6.1 No money: No Shopping and No Participation Cost

To understand the predictions of the model better, we next perform a series of comparative statics. The first model we can compare our results to is the standard portfolio choice model with no stock market participation costs and where money does not circulate in the economy. In our model this specification is nested by setting the shopping technology parameter (ε) equal to zero and the fixed cost of participation (F) also equal to zero. Note that in this case everyone participates in the stock market. Therefore, the composition of stockholders is very different from the baseline model because the households that save very little also participate (aggressively) in the stock market due to the absence of any fixed cost. Figure 6 compares this special case (dotted line) with our baseline estimated model. This is an interesting case because money does not circulate in this economy and the model becomes identical with the recent models on household portfolio choice like Cocco, Gomes, and Maenhout (2005), but with the additional ex ante preference heterogeneity.

We can see that the models that treat money and bonds as perfect substitutes generate a large demand for stocks early in life because future labor income is treated like a bond: all saving is done through the stock market. Moreover, the poorer households also invest fully in the stock market now and therefore the share of wealth in stocks in at one over a very large part of the life cycle. In contrast, the shopping technology model generates a demand for money that produces an upward sloping share of wealth in stocks over some of the early years of the life cycle and also reduces substantially the demand for stocks early in life. The demand for bonds is almost unaffected in the early part of the life cycle but is substantially higher in the baseline model. In the model without any costs, for this particular case of preference heterogeneity, there is a stronger demand for stocks from the poorer households (as a share of their wealth in stocks) and therefore the share of wealth allocated to bonds is lower than in the baseline model almost over the whole part of the life cycle (with the strongest differences during retirement).

The model without money clearly predicts that the share of wealth in stocks decreases as financial wealth rises, a counterfactual prediction. Recently Wachter and Yogo (2010) argue that non-separabilities in the utility function across different goods can generate the upward sloping shape for the share of wealth in stocks as financial wealth increases. Our model provides an alternative explanation that relies on the determinants of money demand and treating transaction accounts and bond investments as assets with different risk/return characteristics.

6.2 High (ten percent) inflation

Figures 7-10 show the results when mean inflation is substantially increased (we set annual inflation equal to ten percent). We present results separately for stockholders and non-stockholders. A high mean inflation decreases the mean rate of return of holding money, and as a result, households reduce money holdings, which is in line with the money demand literature.

Interestingly, Figure 7 shows that inflation affects the portfolios of young and old stockholders differently. Over most parts of the life cycle (except the very beginning when there is very little accumulated financial wealth), the share of wealth in money is lower relative to the baseline model but the reduction is stronger early in life. Young stockholders increase their share of wealth in stocks (up to age 40) by reducing money holdings in the portfolio. Older stockholders, on the other hand, increase bond holdings slightly instead, leaving stocks broadly unchanged.

Figure 8 illustrates how the behavior of non-stockholders is affected by higher inflation. Interestingly, there is a substantially stronger effect in the re-allocation from money into bonds for this group than for the stockholders (as a proportion of total financial wealth). This is intuitive: this group holds a very large fraction of its wealth in money and its money demand is therefore sensitive to the return from holding money relative to bonds.

Figure 9 shows that the higher level of inflation impacts stockholder consumption and wealth accumulation negatively throughout the life cycle. Higher shopping costs and lower returns on money holdings lead to a decline in financial wealth accumulation but this decline is mitigated by the higher mean returns from investing a larger share of the portfolio in stocks, especially in the early part of the life cycle. As a result, financial wealth accumulation is not reduced substantially. Consumption falls slightly because households hold less money due to its lower return and try to limit the rise in shopping costs by spending less, and there is also a wealth effect. Non-stockholders' wealth is affected mainly in youth but due to the low wealth of these agents consumption remains broadly unchanged since real wages are unaffected by higher mean inflation.

Figure 10 shows that higher inflation increases slightly stock market participation throughout life as households hedge higher mean inflation by investing in the stock market.

6.3 Zero inflation

We next examine the implications of a zero inflation regime on saving and portfolio choices. The results are shown in Figures 11-14. We again show first the results for stockholders. In this case, the mean return on money increases relative to the benchmark. As is shown in Figure 11, the average share of wealth allocated to money significantly increases for stockholders aged 30 and higher. Bond holdings decline substantially over this part of the life cycle, as households substitute into money. Average stockholdings increase for households in their 20s and decrease for households in their 30s, with stockholdings broadly unchanged for older age groups. The change in the stock-ownership behavior of young households explains the peculiar response of money holdings for young stockholders until around age 40. This is because the only young households that remain stockholders in their 20s in the low inflation environment are the lucky few that receive large enough bequests at birth and begin participating immediately. Since they are richer than average, they allocate a higher percentage of their financial wealth into stocks. It is this change in the composition of stockholders which explains much of the portfolio shifts observed in Figure 11 early on in the life cycle.

Figure 12 shows how non-stockholders change their behavior in a zero inflation environment. There is a large increase in money demand and a decline in the demand for bonds throughout the life cycle.

Figure 13 gives us more information about the way zero inflation affects consumption and wealth accumulation of different groups. For non-stockholders, there is a slightly higher wealth accumulation (relative to the baseline case) as the cut-off level of financial wealth to participate in the stock market is now higher. This happens because the additional urgency to participate in the stock market to avoid the effects of inflation is missing when mean inflation is zero. For stockholders, consumption and wealth accumulation actually decline slightly over the largest part of the life cycle (after around age 35), but both are higher in the early part of the life cycle. The early part is the composition effect pointed out above: these stockholders are richer and therefore have higher consumption and wealth over this part of the life cycle. Given the higher cut-off wealth point required to enter the stock market, indicated by the lower stock market participation early in life (figure 14), stockholders take advantage of the higher equity premium later in life and/or after a larger wealth accumulation level than before, resulting in lower wealth accumulation and consumption in the later parts of the life cycle.

Figure 14 illustrates that there is a substantial decrease in the stock market participation rate during the deflationary period because money now provides a better rate of return to its holders. This decrease is especially pronounced for younger households that delay their entry into the stockmarket until their 30s (rather than their 20s as in the benchmark).

Interestingly the change in household portfolios is larger in absolute magnitudes compared to the high inflation scenario. This shows that money demand becomes more interest elastic at lower rates of inflation. The model also has the interesting prediction that the impact of inflation on the rate of stockmarket participation should also be larger at low rates of inflation.

6.4 Hedging Demands

We next examine how various volatilities and correlations of shocks affect portfolio choice over life cycle. When mean inflation is low at 2.5% as in the benchmark case, we find that increasing inflation volatility has little effects on portfolio choice. More specifically, increasing the standard deviation of inflation by 5 percentage points, keeping mean inflation at 2.5%, does not have significant effects on household behavior. There is a slight increase in stock market participation over the life cycle and for stockholders a slightly more aggressive tilt towards stocks (and away from money) in the financial portfolio over the whole life cycle but the differences from the baseline model are quantitatively very small.

We also examine the effects of changing the correlation of inflation with bonds and stock returns through two experiments. First, we set the correlation between inflation and stock returns equal to zero, keeping the other correlations the same as in the benchmark case. In the second, we set the correlation of inflation with bond returns equal to zero. Both experiments show that changing these correlations have negligible quantitative effects on financial wealth accumulation and portfolio choices.

We conclude that when mean inflation is low at 2.5%, hedging demands due to inflation volatility are very small. One hypothesis might be that hedging demands due to volatile inflation may be important only when inflation is high. Data show that inflation volatility is positively correlated with the mean inflation rate. For this purpose, we also consider the economy with a mean inflation rate at 10%, and examine the effects of higher inflation volatility in this regime. Comparing the cases in which the standard deviation of inflation is 1% and 5%, respectively, we find very small differences on portfolio choice. In the interest of space we do not report those figures.

What is they intuition behind these results? The intuition is that most risk households face is idiosyncratic and not aggregate. The volatilities of aggregate risks (wages and inflation) are an order of magnitude lower than idiosyncratic labor income risk. Even very high correlations between aggregate variables do not translate to individual hedging demands exactly because they have a relatively small effect on the idiosyncratic risks households face. This points towards building models with stronger persistence in the aggregate variables and/or models with long-run co-integrating relationships between aggregate variables (like, for instance, Benzoni, Collin-Dufresne, and Goldstein (2007)). Such long run co-integrating relationships (for instance between the aggregate component of labor income and dividends) can increase the risk of stock market investments and further reduce the demand for stocks early in life. At the same time they might generate stronger hedging demands. We leave this interesting and important extension for future research.

7 Conclusion

We estimate the preference parameters of a life cycle money demand and portfolio choice model. The predictions of the model are consistent with the data on limited stock market participation, money demand and wealth accumulation over the life cycle but the model overpredicts allocations to the stock market in the early period of life (and therefore underpredicts bond holdings until age 44). We use the model to analyze how inflation or deflation affect money demand, asset allocation and stock market participation over the life cycle. Higher inflation reduces money demand and increases the shares of bonds in household portfolios. For older stockholders and non-stockholders, the main adjustment margin is between money and bonds. For younger stockholders, the main margin is between money and stocks. Stock market participation increases with higher inflation.

Future work can extend the analysis using other preferences (for instance, see Calvet and Sodini (2014)) or with more accurate and detailed administrative data that allow a cleaner separation between different financial assets (for instance, see Fagereng, Gottlieb, and Guiso (2014)) or in a general equilibrium setting to address policy questions such as the effects of changes in the relative supply of money and bonds through open market operations by the central bank. Relatedly, a general equilibrium model can be used to compute the welfare cost of inflation across different households, and in the aggregate.

Appendix A The Data

A.1 Survey of Consumer Finances

We use repeated cross sections from the U.S. Survey of Consumer Finances to establish certain robust facts with regards to household choices across liquid accounts (money), bonds and stocks. Total financial assets are broken up into the three broad categories the model has implications for: liquid resources (LIQ), stock (EQUITY) and nonequity (BOND) investments. In the 2001 public extract of the SCF data set, LIQ is defined as the sum of all checking, saving, money market deposit and call accounts. We follow the same convention and

LIQ becomes our measure of money when confronting the model implications to the data. EQUITY is defined in the same extract as all financial assets invested in stocks and this comprises the following categories:

1) directly held stock

2) stock mutual funds (the full value is assigned if the fund is described as a stock mutual fund, and half the value for combination mutual funds)

3) IRAs/Keoghs invested in stock (full value if mostly invested in stock, half value if split between stocks/bonds or stocks/money market, one third value if split between stocks/bonds/money market),

4) other managed assets with equity interest (annuities, trusts, MIAs) (where again the full value is used if mostly invested in stock, half value if split between stocks/MFs & bonds/CDs, or "mixed/diversified," and one third value if "other")

5) thrift-type retirement accounts invested in stock (full value if mostly invested in stock and half value if split between stocks and interest earning assets) and

6) savings accounts classified as 529 or other accounts that may be invested in stocks. We classify the remaining financial assets as BOND and interpret them as capturing the bond investments in the model (both government and corporate bonds are lumped together in this category).

A.2 Aggregate Data

We use the CRSP data base to download annual US inflation, bond and stock returns from 1925 to 2008. We use the 1995-2008 averages because we focus on matching the data from the 2007 SCF survey but report results for other samples as well. For the aggregate component of labor income we use the National Income and Product Accounts (NIPA) wages and salary disbursement series and we deflate using the inflation rate from the Center of Research in Security Prices (CRSP).

Appendix B The Normalized Value Function

Let $v_{it}^j \equiv V_{it}/Y_{it}^p$ be the normalized value of individual i at age t. Households also differ according to their preferences, denoted by $j = H, L$. Households of type H have a high risk aversion (ρ_H), high EIS (ψ_H) and a high bequest motive (φ_H). Households of type L have low risk aversion (ρ_L), low EIS (ψ_L) and no bequest motive ($\varphi_L = 0$). $g_{it+1} \equiv Y_{it+1}^p/Y_{it}^p$ is the growth rate of the permanent component of income for the household.

$$
\begin{aligned}
v_{it}^j(x_{it}, I_t = a) &= \Bigg[(1-\beta)c_{it}^{1-1/\psi_j} \\
&\quad + \beta \left\{ E_t \left[\begin{array}{c} \xi_t (v_{it+1}^j (x_{it+1}, I_{t+1} = a))^{1-\rho_j} (Y_{it+1}^p/Y_{it}^p)^{1-\rho_j} + \\ (1-\xi_t)\varphi_j (x_{it+1})^{1-\rho_j} (Y_{it+1}^p/Y_{it}^p)^{1-\rho_j} \end{array} \right] \right\}^{\frac{1-1/\psi_j}{1-\rho_j}} \Bigg]^{\frac{1}{1-1/\psi_j}} \\
&= \left[\begin{array}{c} (1-\beta)c_{it}^{1-1/\psi_j} + \\ \beta \left\{ E_t \left[\begin{array}{c} \xi_t (v_{it+1} (x_{it+1}, I_{t+1} = a)\, g_{it+1})^{1-\rho_j} + \\ (1-\xi_t)\varphi_j (x_{it+1} g_{it+1})^{1-\rho_j} \end{array} \right] \right\}^{\frac{1-1/\psi_j}{1-\rho_j}} \end{array} \right]^{\frac{1}{1-1/\psi_j}},
\end{aligned}
\tag{B.1}
$$

for $a = 0, 1$. The continuous state is x_{it} (normalized cash on hand) and its evolution is given by (8). x_{it} is the cash on hand which is at the disposal of the household before the payment of the stock market participation cost F (this is important later). The state also includes participation status (denoted by I_t) where 1 denotes participation and 0 denotes non-participation.

Appendix C Numerical Solution

We exploit the scale-independence of the maximization problem and rewrite all variables as ratios to the permanent component of labor income (Y_{it}^p). The laws of motion and the value function can then be rewritten in terms of these normalized variables, and we use lower case letters to denote them. This normalization allows us to reduce the number of state variables to three: liquid wealth, participation status, and age. The problem is solved as follows.

For households who are already stock market participants, there is no participation decision. Their value function is given by:

$$v_t^j(x_{it}, I_t = 1) = \underset{c_t, \alpha_t^s, \alpha_t^b}{MAX} \left\{ (1-\beta) c_t^{1-1/\psi_j} \right.$$
$$\left. + \beta \left(E_t \left\{ \left(\frac{Y_{it+1}^p}{Y_{it}^p} \right)^{1-\rho_j} \left(\begin{array}{c} \xi_t [v_{t+1}^j(x_{it+1}, I_{t+1} = 1)]^{1-\rho_j} \\ +(1-\xi_t)\varphi_j (x_{it+1})^{1-\rho_j} \end{array} \right) \right\} \right)^{\frac{1-1/\psi_j}{1-\rho_j}} \right\}^{\frac{1}{1-1/\psi_j}}$$

Non-participants decide whether or not to incur the fixed cost F at time (age) t and this immediately comes out of their cash on hand x_{it}. They compare the two value functions associated with direct stock market participation or continued non-participation:

$$v_t^j(x_{it}, I_t = 0) = MAX_{0,1}\{v_t^j(x_{it}, I_t = 0), v_t^j(x_{it} - F, I_t = 1)\}$$

where $I_t = 1$ denotes stock market participation. The value of remaining a non-participant is given by:

$$v_t^j(x_{it}, I_t = 0) = \underset{c_t,\alpha_t^b}{MAX}\left\{(1-\beta)c_t^{1-1/\psi_j} + \beta\left(E_t\left\{\left(\frac{Y_{it+1}^p}{Y_{it}^p}\right)^{1-\rho_j}\left(\begin{array}{c}\xi_t[v_{t+1}^j(x_{it+1}^0, I_{t+1}=0)]^{1-\rho_j}\\ +(1-\xi_t)\varphi_j\left(x_{it+1}^0\right)^{1-\rho_j}\end{array}\right)\right\}\right)^{\frac{1-1/\psi_j}{1-\rho_j}}\right\}^{\frac{1}{1-1/\psi_j}}$$

where

$$x_{it+1}^0 = \frac{r_{t+1}^b}{g_{it+1}}b_{it} + \frac{r_{t+1}^m}{g_{it+1}}m_{it} + U_{it+1} - \frac{\omega_{it}}{g_{it+1}}$$

is normalized cash-on-hand in period $t+1$ conditional on the decision not to begin stock market participation at time t.

For those who decide to participate

$$v_t^j(x_{it} - F, I_t = 1) = \underset{c_t,\alpha_t^s,\alpha_t^b}{MAX}\left\{(1-\beta)c_t^{1-1/\psi_j} + \beta\left(E_t\left\{\left(\frac{Y_{it+1}^p}{Y_{it}^p}\right)^{1-\rho_j}\left(\begin{array}{c}\xi_t[v_{t+1}^j(x_{it+1}^1, I_{t+1}=1)]^{1-\rho_j}\\ +(1-\xi_t)\varphi_j\left(x_{it+1}^1\right)^{1-\rho_j}\end{array}\right)\right\}\right)^{\frac{1-1/\psi_j}{1-\rho_j}}\right\}^{\frac{1}{1-1/\psi_j}}$$

where

$$x_{it+1}^1 = \frac{r_{t+1}^s}{g_{it+1}}s_{it} + \frac{r_{t+1}^b}{g_{it+1}}b_{it} + \frac{r_{t+1}^m}{g_{it+1}}m_{it} + U_{it+1} - \frac{\omega_{it}}{g_{it+1}}$$

is normalized cash-on-hand in period $t+1$ conditional on the decision to begin stock market participation at time t. The main difference between x_{it+1}^0 and x_{it+1}^1 lies in the fact that x_{it+1}^1 includes returns from holding stocks ($\frac{r_{t+1}^s}{g_{it+1}}s_{it}$).

We solve the model recursively backwards starting from the last period for households of each type $j = H, L$. In the last period ($t = T$) the policy functions are trivial and the value function corresponds to the bequest function.

We need to solve for four control variables in every year for stock-holders: current consumption (c_t), the fraction of the portfolio allocated to stocks (α_t^s) and bonds (α_t^b) (the fraction of saving allocated to money α_t^m can be determined as the residual) and the participation decision.

For every age t prior to T, and for each point in the state space, we optimize using grid search. From the Bellman equation the optimal decisions are given as current utility plus the discounted expected continuation value ($E_t v_{t+1}^j(.)$), which we can compute since we have just obtained v_{t+1}^j. We perform all numerical integrations using Gaussian quadrature to approximate the distributions of the innovations to the labor income process and the risky asset returns. Cubic splines are used to perform the interpolation of the value function for points which do not lie on the state space grid, with more points used at lower levels of wealth where the value function has high curvature. Once we have computed the value of each alternative we pick the maximum, thus obtaining the policy rules for the current period. Substituting these decision rules in the Bellman equation, we obtain this period's value function ($v_t^j(.)$), which is then used to solve the previous period's maximization problem. This process is iterated until $t = 1$.

References

AGUIAR, M., AND E. HURST (2013): "Deconstructing Life Cycle Expenditure," *Journal of Political Economy*, 121, 437–492.

ALAN, S. (2006): "Entry Costs and Stock Market Participation over the Life Cycle," *Review of Economic Dynamics*, pp. 588–611.

ALVAREZ, F., L. GUISO, AND F. LIPPI (2012): "Durable consumption and asset management with transaction and observation costs," *American Economic Review*, 102 (5), 2272–2300.

ALVAREZ, F., AND F. LIPPI (2009): "Financial Innovation and the Transactions Demand for Cash," *Econometrica*, 77(2), 363–402.

ATTANASIO, O., J. BANKS, C. MEGHIR, AND G. WEBER (1999): "Humps and bumps in lifetime consumption," *Journal of Business & Economic Statistics*, 17(1), 22–35.

ATTANASIO, O., AND M. PAIELLA (2010): "Intertemporal Consumption Choices, Transaction Costs and Limited Participation in Financial Markets: Reconciling Data and Theory," *Journal of Applied Econometrics*, pp. 322–343.

BAUMOL, W. J. (1952): "The transactions demand for cash: An inventory theoretic approach," *The Quarterly Journal of Economics*, 66(4), 545–556.

BENZONI, L., P. COLLIN-DUFRESNE, AND R. S. GOLDSTEIN (2007): "Portfolio Choice over the Life-Cycle When the Stock and Labor Markets are Cointegrated," *Journal of Finance*, pp. 2123–2167.

BONAPARTE, Y., R. COOPER, AND G. ZHU (2012): "Consumption Smoothing and Portfolio Rebalancing: The Effects of Adjustment Costs," *Journal of Monetary Economics*, 59 (8), 751–768.

CAGETTI, M. (2003): "Wealth Accumulation over the Life Cycle and Precautionary Savings.," *Journal of Business & Economic Statistics*, 21(3), 339–354.

CALVET, L., AND P. SODINI (2014): "Twin Peaks: Disentangling the Determinants of Risk-Taking in Household Portfolios," *Journal of Finance*, (2).

CAMPBELL, J. Y. (2006): "Household Finance," *Journal of Finance*, 61(4), 1553–1604.

CARROLL, C. D. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112(1), 1–55.

Castaneda, A., J. Diaz-Gimenez, and J.-V. Rios-Rull (2003): "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, 111(4), 818–857.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005): "Consumption and portfolio choice over the life cycle," *Review of financial Studies*, 18(2), 491–533.

Cooper, R., and G. Zhu (2013): "Household Finance: Education, Permanent Income and Portfolio Choice," *Penn State working paper*, (2).

De Nardi, M. (2004): "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71, 743–768.

De Nardi, M., E. French, and J. Jones (2010): "Differential mortality, uncertain medical expenses, and the saving of elderly singles," *Journal of Political Economy*, pp. 39–75.

Deaton, A. (1991): "Saving and liquidity constraints," *Econometrica*, 59(5), 1221–1248.

Doepke, M., and M. Schneider (2006): "Inflation and the Redistribution of Nominal Wealth," *Journal of Political Economy*, 114(6), 1069–1097.

Duffie, D., and K. Singleton (1993): "Simulated moments estimation of Markov models of asset prices," *Econometrica*, pp. 929–952.

Epstein, L. G., and S. E. Zin (1989): "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica: Journal of the Econometric Society*, pp. 937–969.

Fagereng, A., C. Gottlieb, and L. Guiso (2014): "Asset Market Participation and Portfolio Choice over the Life-Cycle," *Working paper, Ente Einaudi.*

Gomes, F., and A. Michaelides (2005): "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," *Journal of Finance*, 60(2), 869–904.

Gourinchas, P.-O., and J. A. Parker (2002): "Consumption Over the Life Cycle," *Econometrica*, 70(1), 47–89.

Guiso, L., P. Sapienza, and L. Zingales (2008): "Trusting the Stock Market," *The Journal of Finance*, 63, 2557–2600.

Haliassos, M., and A. Michaelides (2003): "Portfolio Choice and Liquidity Constraints," *International Economic Review*, 44(1), 143–177.

Heaton, J., and D. Lucas (2000a): "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *Journal of Finance*, 55(3), 1163–1198.

——— (2000b): "Portfolio Choice in the Presence of Background Risk," *Economic Journal*, pp. 1–26.

Kiyotaki, N., and R. Wright (1989): "On Money as a Medium of Exchange," *Journal of Political Economy*, 97(4), 927–54.

Krusell, P., and A. A. J. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867–896.

Lippi, F., and A. Secchi (2009): "Technological change and the hhousehold' demand for currency," *Journal of Monetary Economics*, 56,2, 222–230.

Ljungqvist, L., and T. J. Sargent (2004): *Recursive macroeconomic theory.* The MIT Press.

Love, D. (2010): "The Effects of Marital Status and Children on Savings and Portfolio Choice," *Review of Financial Studies*, pp. 385–432.

Lucas, R. E. (2000): "Inflation and Welfare," *Econometrica*, 68(2), 247–274.

Lucas, R. E., and N. Stokey (1987): "Money and Interest in a Cash-in-Advance Economy," *Econometrica*, 55(3), 491–513.

McCallum, B. T., and M. S. Goodfriend (1987): "Demand for money: theoretical studies," *The New Palgrave: A Dictionary of Economics*, 1, 775–81.

Mulligan, C. B., and X. Sala-i-Martin (2000): "Extensive margins and the demand for money at low interest rates," *Journal of Political Economy*, 108(5), 961–991.

Paiella, M. (2007): "The Foregone Gains of Incomplete Portfolios," *The Review of Financial Studies*, pp. 1623–1646.

Polkovnichenko, V. (2007): "Life-Cycle Portfolio Choice with Additive Habit Formation Preferences and Uninsurable Labor Income Risk," *Review of Financial Studies*, 20(1), 83–124.

Sidrauski, M. (1967): "Rational choice and patterns of growth in a monetary economy," *The American Economic Review*, 57(2), 534–544.

Tobin, J. (1956): "The interest-elasticity of transactions demand for cash," *The Review of Economics and Statistics*, 38(3), 241–247.

Vissing-Jorgensen, A. (2002): "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*, 110(4), 825–853.

Wachter, J., and M. Yogo (2010): “Why Do Household Portfolio Shares Rise in Wealth?,” *Review of Financial Studies*, pp. 3929–3965.

Weil, P. (1990): “Nonexpected Utility in Macroeconomics,” *The Quarterly Journal of Economics*, 105(1), 29–42.

Figure 1: Baseline policy functions for Low Elasticity of Intertemporal Substitution (EIS) and Low Relative Risk Aversion (RRA) households
(Solid line: stockholders, Dashed line: non-stockholders)

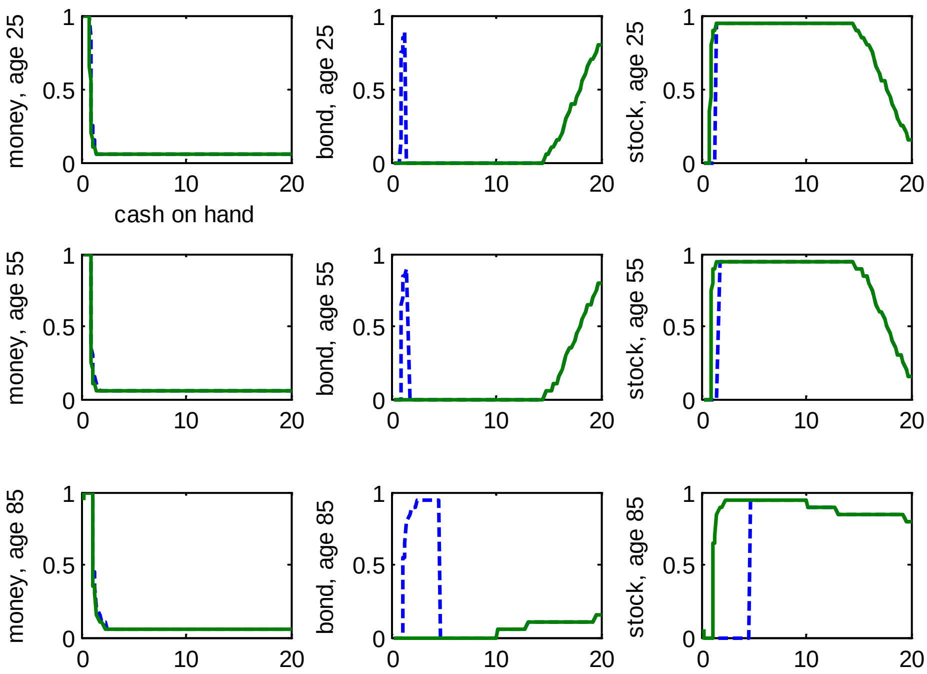

Figure 2: Baseline policy functions for High Elasticity of Intertemporal Substitution (EIS) and High Relative Risk Aversion (RRA) households
(Solid line: stockholders, Dashed line: non-stockholders)

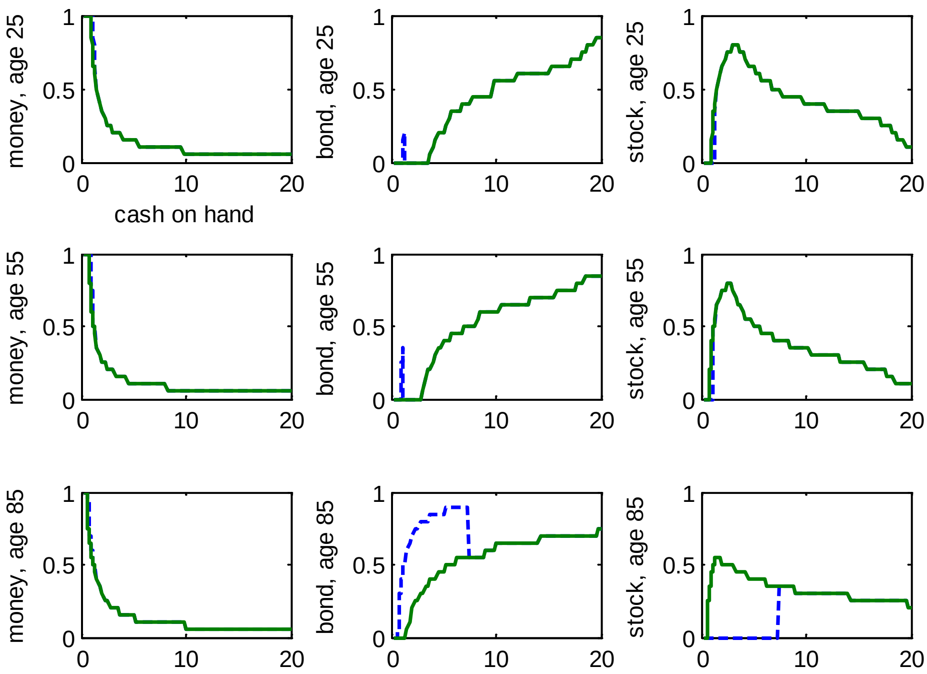

Figure 3: Consumption, financial wealth and income over the life cycle: Baseline model

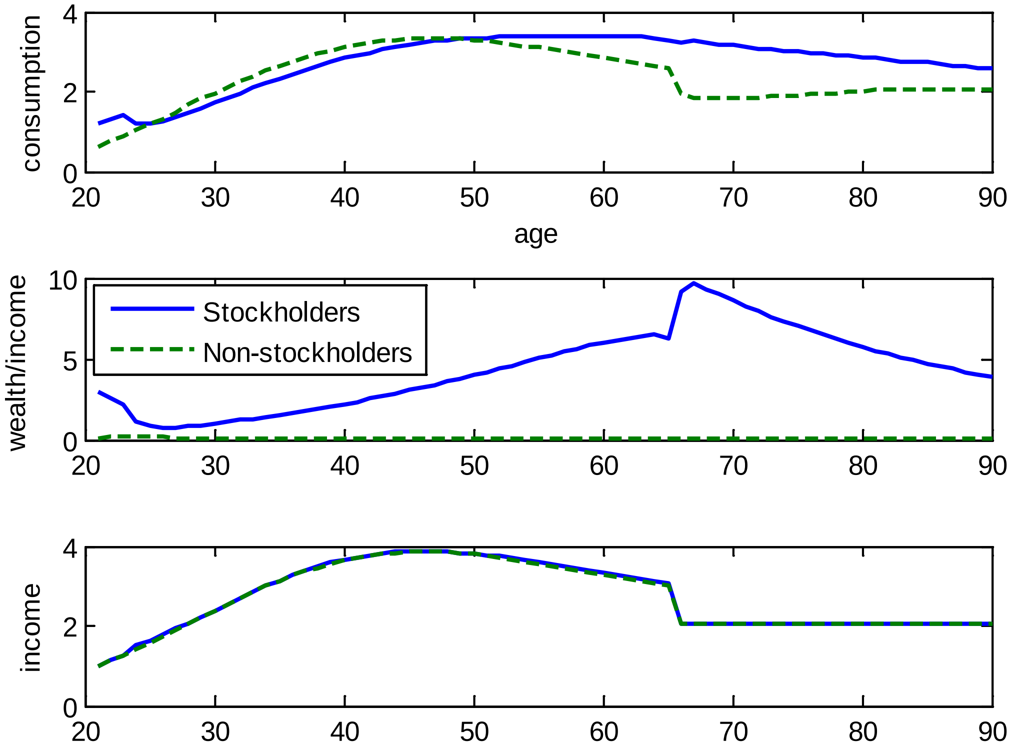

Figure 4: Portfolio choice over the life cycle: Baseline model

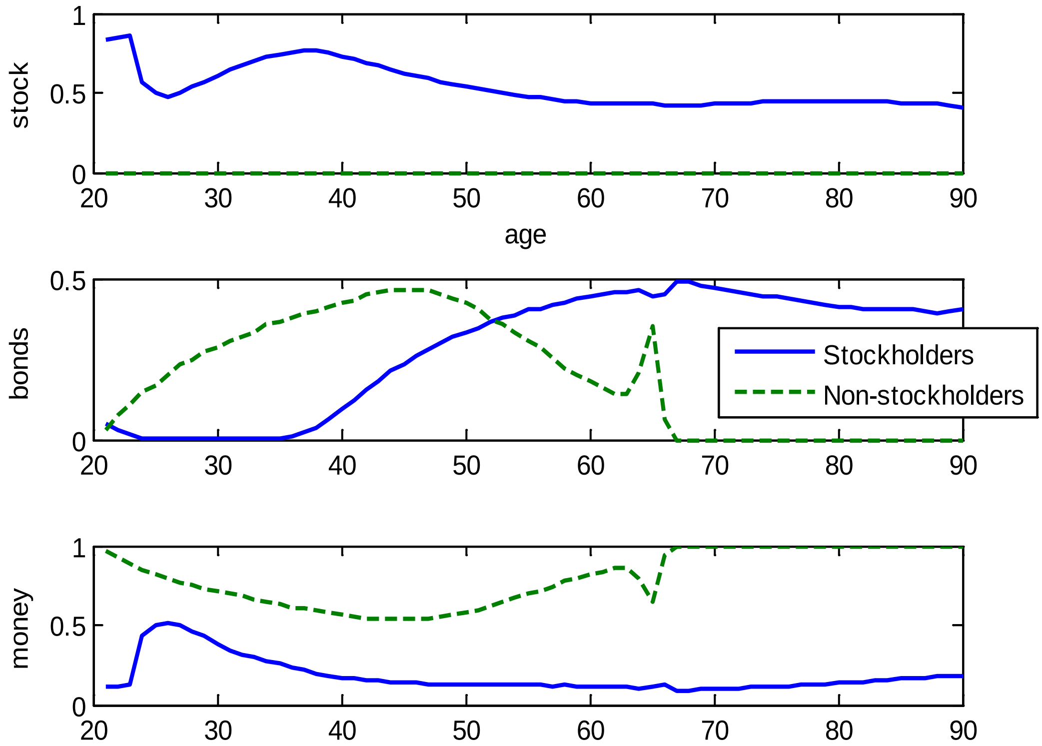

Figure 5: Stock market participation over the life cycle: Baseline model

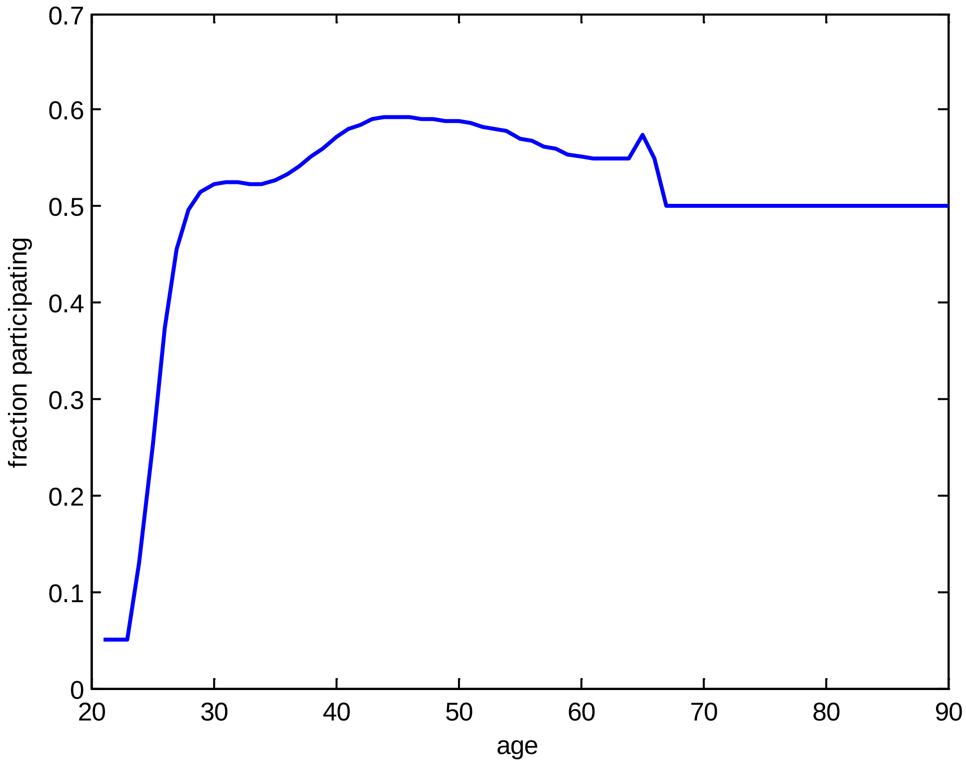

Figure 6: Stockholder portfolios over the life-cycle: No shopping and participation costs versus Baseline Model

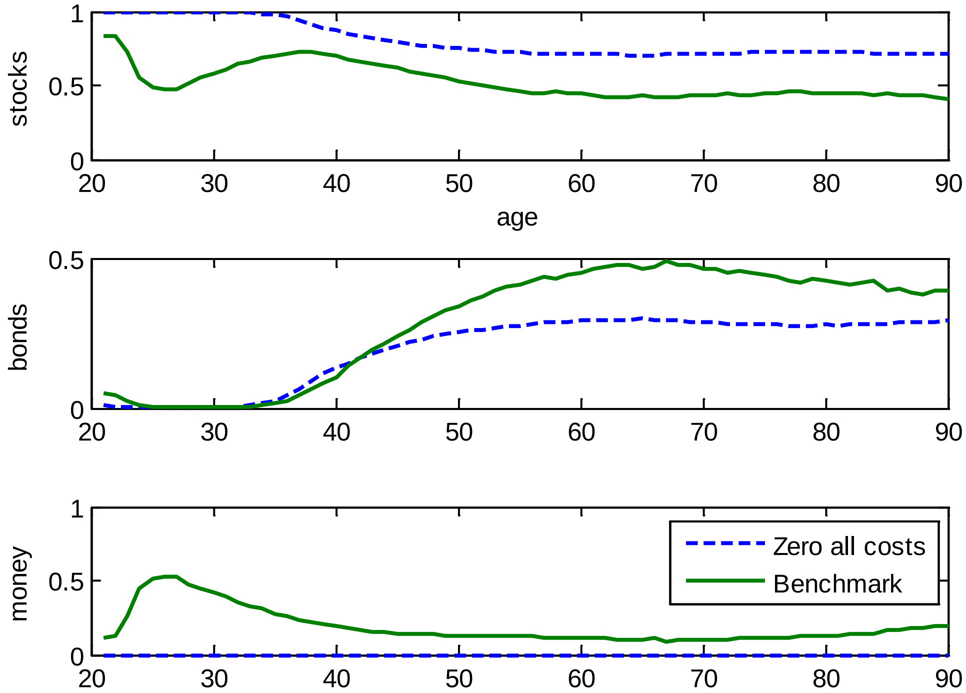

Figure 7: Stockholder portfolios over the life cycle: High Mean Inflation (10% per annum) versus Baseline Model (2.5% per annum)

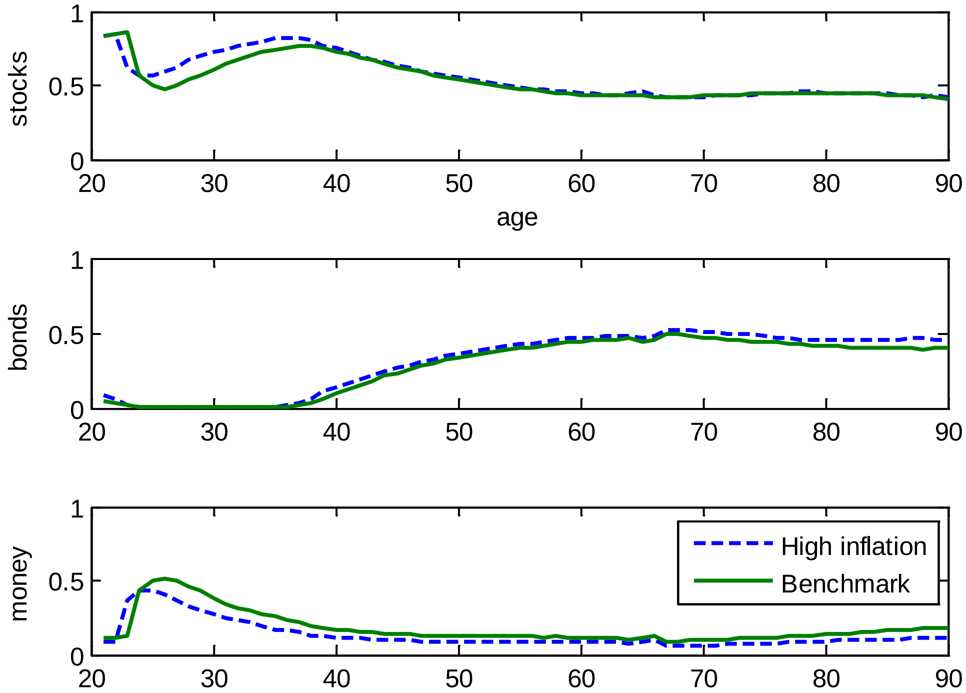

Figure 8: Non-stockholder portfolios over the life cycle: High Mean Inflation (10% per annum) versus Baseline Model (2.5% per annum)

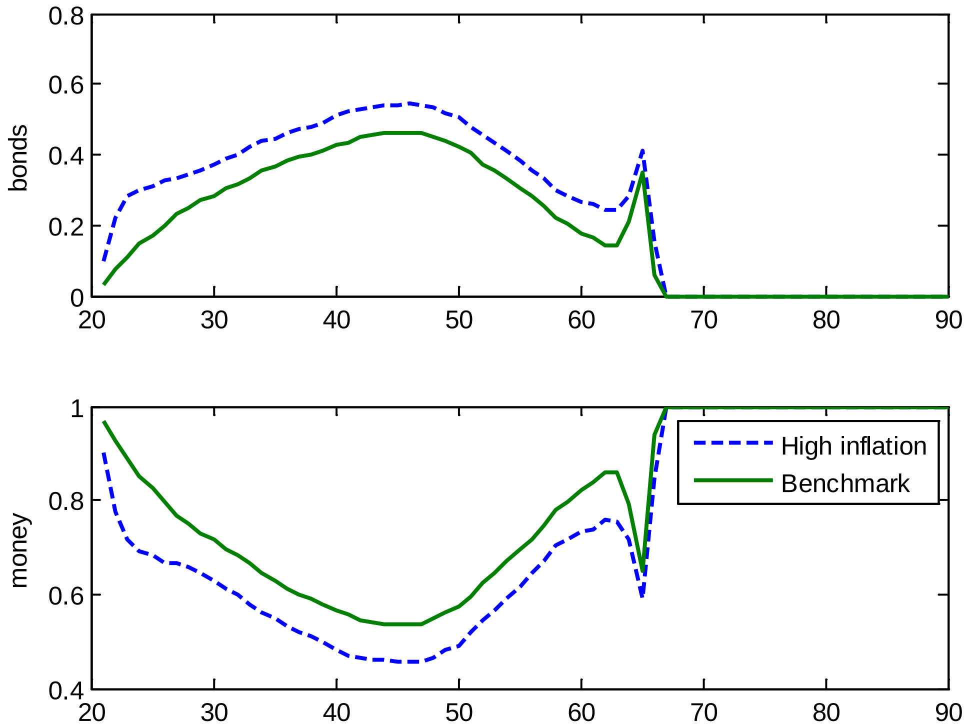

Figure 9: Consumption, financial wealth and income over the life cycle: High Mean Inflation (10% per annum) versus Baseline Model (2.5% per annum)

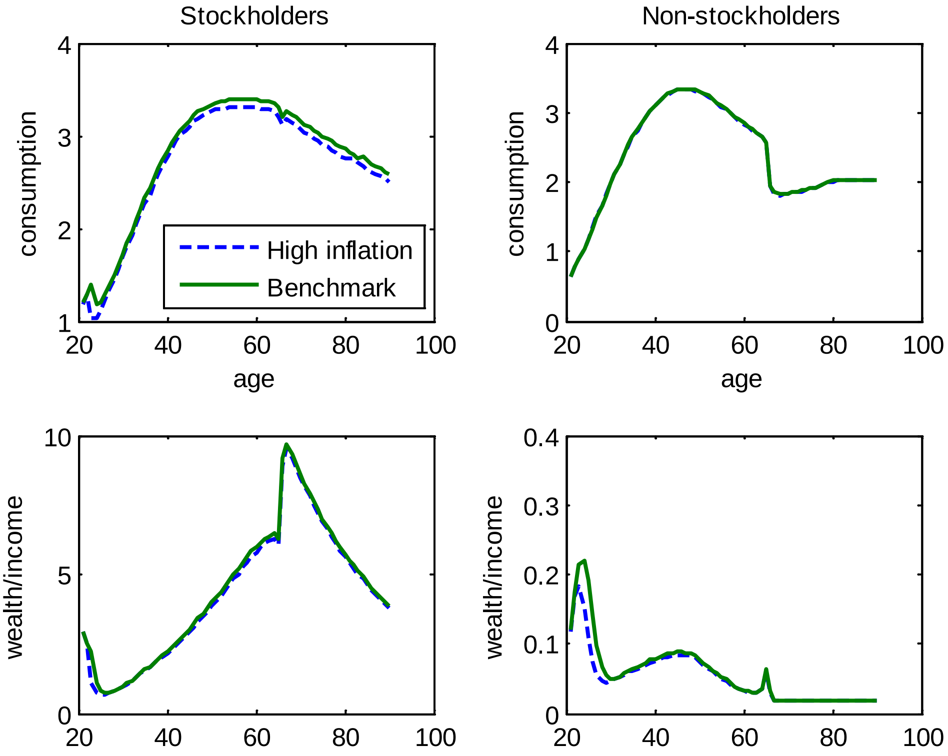

Figure 10: Stock market participation: High Mean Inflation (10% per annum) versus Baseline Model (2.5% per annum)

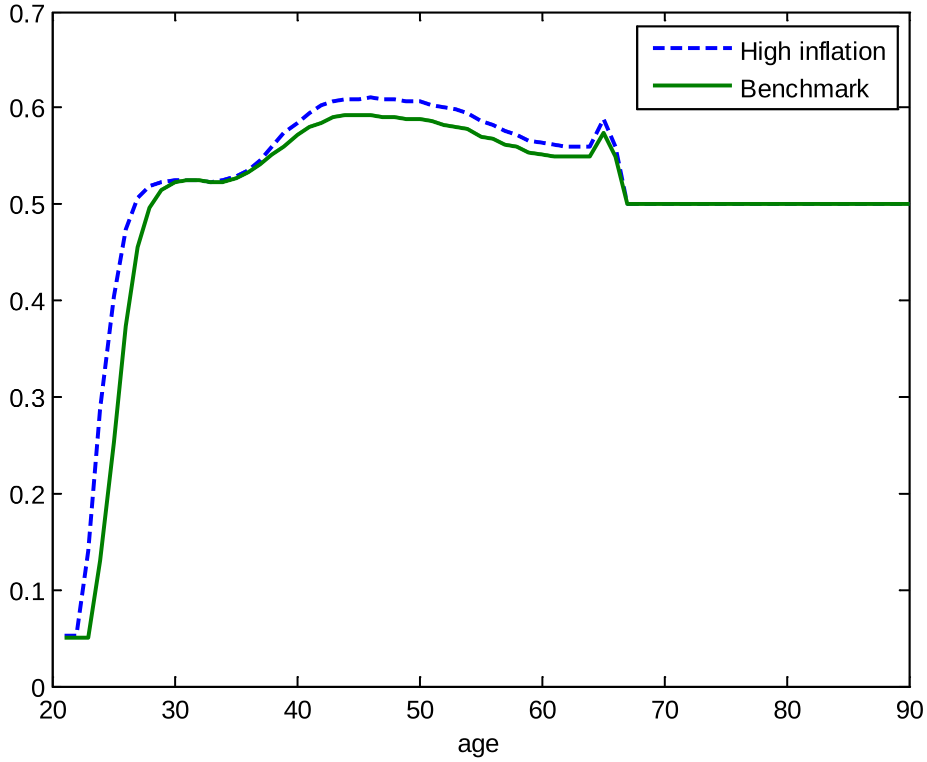

Figure 11: Stockholder Portfolios over the life cycle: Zero Mean Inflation versus Baseline Model (2.5% per annum)

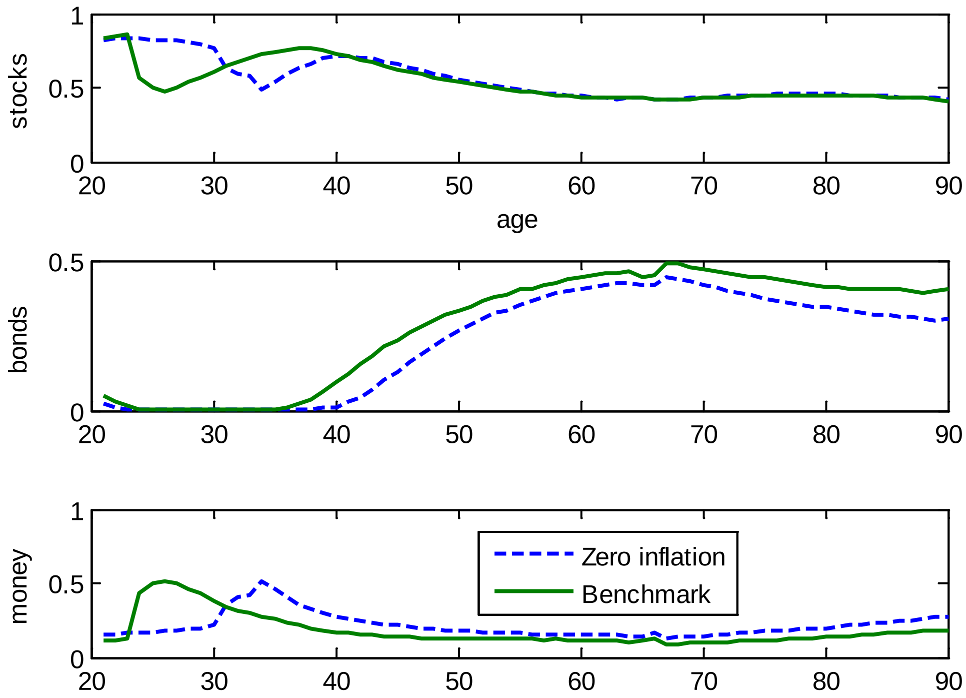

Figure 12: Non-stockholder portfolios over the life cycle: Zero Mean Inflation versus Baseline Model (2.5% per annum)

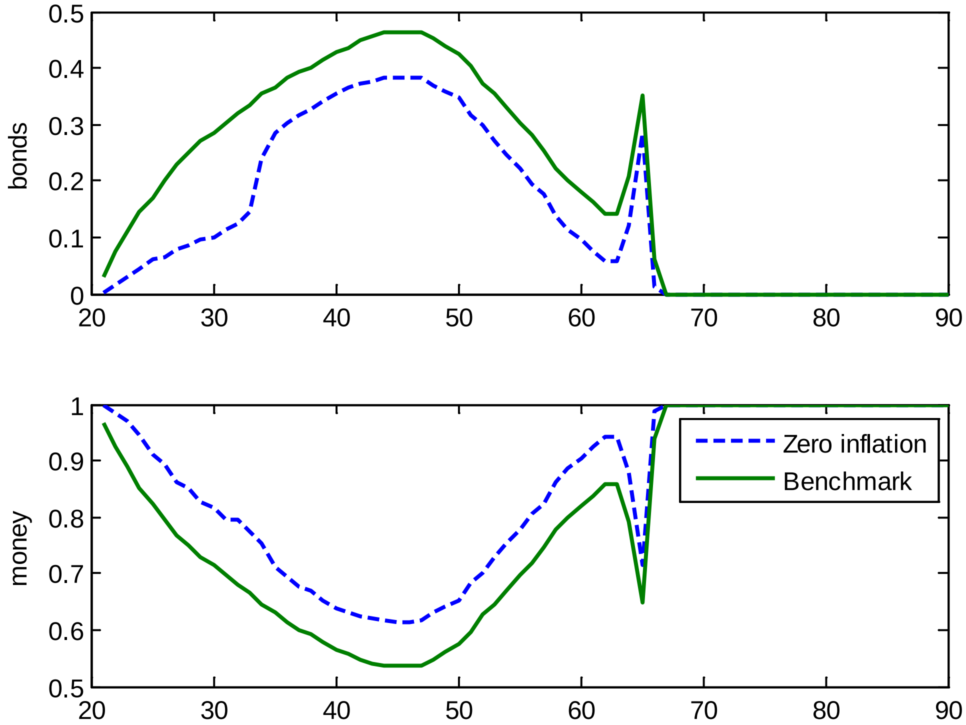

Figure 13: Consumption, financial wealth and income over the life cycle: Zero Mean Inflation versus Baseline Model (2.5% per annum)

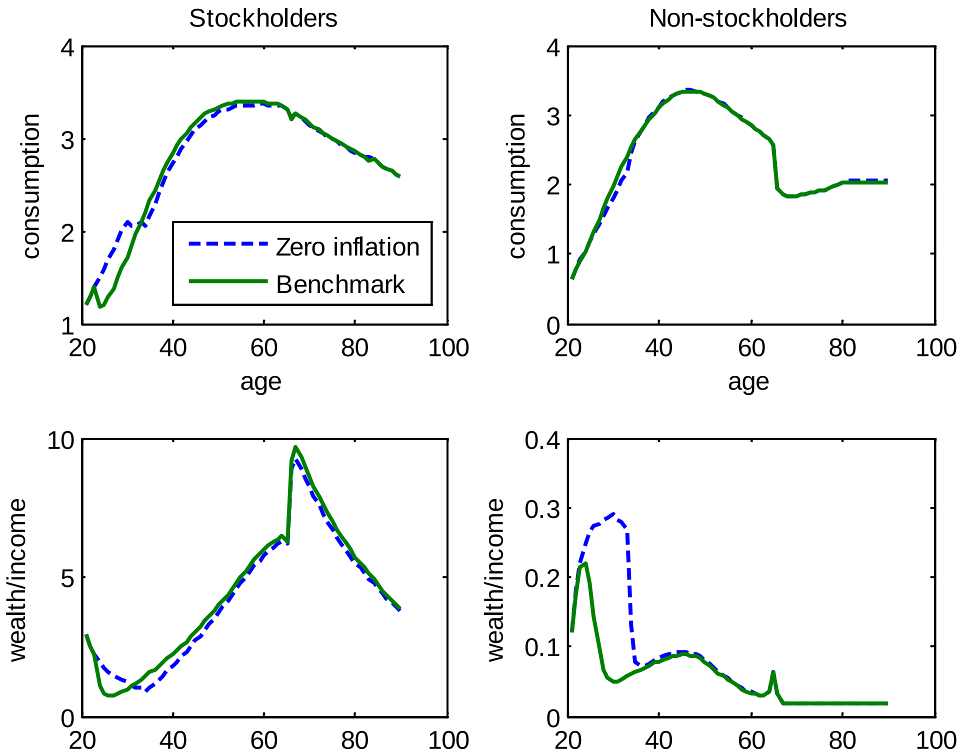

Figure 14: Stock Market Participation: Zero Mean Inflation versus Baseline Model (2.5% per annum)