

DISCUSSION PAPER SERIES

No. 10369

LIFE-CYCLE PORTFOLIO CHOICE WITH LIQUID AND ILLIQUID ASSETS

Claudio Campanale, Carolina Fugazza
and Francisco J. Gomes

FINANCIAL ECONOMICS



Centre for Economic Policy Research

LIFE-CYCLE PORTFOLIO CHOICE WITH LIQUID AND ILLIQUID ASSETS

Claudio Campanale, Carolina Fugazza and Francisco J. Gomes

Discussion Paper No. 10369

January 2015

Submitted 16 January 2015

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **FINANCIAL ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Claudio Campanale, Carolina Fugazza and Francisco J. Gomes

LIFE-CYCLE PORTFOLIO CHOICE WITH LIQUID AND ILLIQUID ASSETS[†]

Abstract

Traditionally, quantitative models that have studied households' portfolio choices have focused exclusively on the different risk properties of alternative financial assets. We introduce differences in liquidity across assets in the standard life-cycle model of portfolio choice. More precisely, in our model, stocks are subject to transaction costs, as considered in recent macro literature. We show that, when these costs are calibrated to match the observed infrequency of households' trading, the model is able to generate patterns of portfolio stock allocation over age and wealth that are constant or moderately increasing, thus more in line with the existing empirical evidence.

JEL Classification: D91, G11 and H55

Keywords: cash-in-advance, household portfolio choice, self-insurance and transaction cost

Claudio Campanale claudio@ua.es
Universidad de Alicante

Carolina Fugazza fugazza@cerp.unito.it
CeRP

Francisco J. Gomes FGomes@london.edu
London Business School and CEPR

[†] We would like to thank Urban Jermann (The Editor) and an anonymous referee for many precious comments and suggestions. We also want to thank Rui Albuquerque, Giovanna Nicodano, seminar participants at Collegio Carlo Alberto, École Polytechnique, IHS, Universidad de Alicante, Universidad Carlos III and Università di Bologna, and conference attendants at the Netspar International Pension Workshop 2011, Midwest Macroeconomics Meetings 2012, SNDE annual symposium 2013 for helpful comments and suggestions. Claudio Campanale wishes to thank the Ministerio de Educación y Ciencia proyecto SEJ 2007-62656, proyecto Prometeo/2013/037 and Collegio Carlo Alberto for financial support and CeRP for generous hospitality during the development of this project. Any remaining errors or inconsistencies are entirely our responsibility.

1 Introduction

The last decade has witnessed a substantial surge of academic interest in the problem of households' financial decisions. A number of empirical facts have been documented regarding in particular the stockholding behavior of households. These include the moderate (albeit increasing) stock market participation rates and the equally modest share allocated to stocks by those who do participate in the stock market. It has also been documented that the share of financial wealth allocated to stocks is increasing in wealth and roughly constant or moderately increasing in age.¹ Equally important has been the development of life-cycle models of portfolio choice that incorporate frictions, constraints, and key sources of risk. These models generate a puzzle that is the extensive-margin equivalent of the equity premium puzzle: given the historical equity premium, households should invest most of their financial wealth in stocks, something that is at odds with the empirical evidence. In the context of asset allocation decisions this puzzle is further compounded with the fact that the patterns of stock holdings by wealth and age are also

¹Among the papers that have uncovered the patterns of household financial behavior are Ameriks and Zeldes (2004), Bertaut and Starr-McCluer (2000) and Heaton and Lucas (2000) for the US. The book by Guiso et al. (2001) documented the same facts for a number of other industrialized countries as well and the work by Calvet et al.(2007) has gone in much greater details to document stock-holding behavior among Swedish households.

inconsistent with the data.

The current paper adds to this latter line of research by exploring the role played by differences in the liquidity of different classes of financial assets. In order to do this we essentially augment the standard life-cycle model of Cocco et al.(2005) with the monetary model in Alvarez et al. (2002). More precisely we assume that agents receive a stochastic uninsurable earnings stream during working life and face both borrowing and no short sale constraints. They have access to two assets, one riskless and one risky (equities). As in Alvarez et al. (2002) we assume that the assets are held in separate accounts, respectively stock account and monetary/liquid account, and that transactions between these two accounts require payment of a fixed cost.

Households receive their wages in the monetary account and a cash-in-advance constraint holds, so that consumption goods can only be purchased with the available money. This gives the liquid asset an advantage as an asset to insure consumption levels early in life, and this advantage is stronger the greater the transaction cost. Similarly a retired agent who is using accumulated wealth to supplement her pension income would like to hold a certain balance in the liquid account rather than paying the fixed cost in every period. In the paper, and following the literature, we model this as a pure monetary cost, but it is also meant to capture the time and information processing cost that is involved in making the associated financial plan. This

cost is then reflected in the frequency of transactions that we observe among households.^{2,3}

The standard model with no transaction costs can only generate the well known policy functions for the stock share that start at 100 percent when the agent has very little wealth and then monotonically decline as wealth increases.⁴ In the model presented here the current share of stocks becomes a state variable. The optimal stock share decision depends on the current stock share — as well as current wealth and earnings — and displays more complex shapes that include patterns that are increasing in wealth especially when both wealth and current earnings are small. The model then generates a life-cycle stock share profile that is either hump-shaped or moderately increasing, depending on the parametrization used. With respect to wealth the simulated data show portfolio allocations to stocks that are increasing

²The empirical evidence in this respect shows that transactions in stock accounts are rare for a large fraction of households, suggesting that once the planning costs are factored in the overall cost is non-trivial (see Biliias et al. (2010) and the Investment Company Institute report “Equity Ownership in America”(2005)).

³An alternative approach is to assume observation costs (e.g. Abel et al. (2007)). Alvarez et al. (2012) construct a model with both observation and transaction costs, and find stronger empirical support for the latter. This lends support to our choice to study the behavior of conditional portfolio shares under infrequent portfolio adjustment by assuming a fixed transaction rather than an observation cost.

⁴This holds under the assumption of no or small correlation between earnings and risky returns. More discussion on this issue will be given later.

over the bottom to mid quartiles of the distribution and then level off or moderately decline at the top. This occurs also when the behavior of stock shares over wealth is conditioned on age. While still not a perfect match with the data these patterns represent a significant improvement over those produced by conventional models.

Our paper belongs to the growing literature on life-cycle asset allocation with labor income risk.⁵ Particularly related are the recent papers by Benzoni et al. (2007), Gomes and Michaelides (2003), Lynch and Tan (2011), Polkovnichenko (2007) and Wachter and Yogo (2010) which have looked for explanations of patterns of household stock market investment over the life-cycle and over wealth levels. Benzoni et al. (2007) and Lynch and Tan (2011) consider alternative specifications of the labor income process which can also deliver portfolio shares that are increasing in wealth, conditional on age. However, in Benzoni et al. (2007) this effect only takes place early in life, since it is driven by the low-frequency correlation between stock return and labor income. Naturally, as the agent approaches retirement this correlation becomes irrelevant. The objective of their paper is to match the unconditional share as a function of age, so it is only necessary to generate this

⁵As initially explored by Heaton and Lucas (1997 and 2000) and Haliassos and Michaelides (2003) in an infinite horizon setting and by Campbell et al. (2001), Cocco et al. (2005) and Gomes and Michaelides (2005) in a life-cycle setting.

effect early in life. Likewise, in Lynch and Tan (2011) the result is driven by business cycle fluctuations in the conditional distribution of income shocks, and therefore the effect is again only present for young households. Gomes and Michaelides (2003) and Polkovnichenko (2007) generate this increasing pattern by assuming habit formation preferences; however they point out that, in order to get strong effects within this model, the importance of the habit must be very high, and therefore it implies counter-factually high levels of wealth accumulation. Wachter and Yogo (2010) achieve the same result assuming multiple goods, and their model generates an increasing relationship between wealth and the portfolio share of risky assets conditional on age. However, in their preferred calibration, the average life-cycle profile is declining, [and] hence does not match the data very well. We see our theory as complementary to the ones mentioned above. The advantage of our approach is that it allows us to match the weakly increasing pattern of the portfolio share both over the life-cycle and over wealth, conditional on age without the need to resort to any form of correlation between labor earnings and market returns, something that is absent during retirement and is likely to be weak at the end of the working life.

A second related strand of literature includes models of monetary economics that assume a portfolio choice between money and other assets, like capital or bonds, and some frictions. Examples are the papers by Alvarez

et al. (2002), Akyol (2004) and Khan and Thomas (2011). Alvarez et al. (2002) construct a model that is similar to the current one in the assumption about the cash-in-advance constraint on consumption purchases; their model is focused on studying the effects of money injections on interest rates and exchange rates. Their framework though is different from the incomplete market model used here. Akyol (2004) uses the incomplete market model to study the optimality of the Friedman rule when agents have access to two assets, money and a bond. In his model a friction is introduced by assuming that trading in the bond market can be performed only before the uncertainty about labor earnings is resolved. Khan and Thomas (2011) consider a model with endogenous market segmentation and show that it can generate sluggish and persistent adjustments of prices and interest rates to a monetary shock in an endowment economy as well as a hump shaped response of employment and output to productivity shocks.

There is also a growing body of literature in finance that studies the role of inaction in household behavior in assets markets. For example, Chien et al. (2012) show that a model with a small fraction of households that re-balance their portfolio in every period and a large fraction of infrequent traders improves substantially the ability of the theory to explain the large counter-cyclical volatility of aggregate risk compensation. Our model generates such infrequent portfolio adjustments by assuming a fixed transaction cost. In

the literature about household portfolio choice, transaction costs have been traditionally considered on housing transactions (e.g. Cocco (2005) and Yao and Zhang (2005)). Ang et al. (2011) study the portfolio holdings in a model with two risky assets, where one is tradable only at random periods (meant to represent private equity), but their framework and object of study are very different. In this literature the closest paper is Bonaparte et al. (2012), who consider a very similar model with adjustment costs of portfolio rebalancing. The focus of the paper is however very different. They estimate the model to match the portfolio rebalancing behavior of investors, while we are particularly interested in the conditional equity allocations as a function of wealth.

Finally this research is also related to recent papers that have tried to estimate the relationship between wealth changes and the share invested in risky assets using a panel data approach on individual household data. These include the works of Brunnermeier and Nagel (2008) and Chiappori and Paiella (2011). These papers find only a weak relationship between wealth and households' risky investment. The current paper, by generating a non-monotone relationship between the stock share of market participants and wealth, may help rationalize those findings.

The rest of the paper is organized as follows. In section 2 we present the description of the model, in section 3 we report the choice of parameters, in

section 4 we report the main findings of the analysis and finally in section 5 some short conclusions are outlined. The paper concludes with an appendix providing a discussion of the numerical methods and a description of the data construction.

2 The Model

The model is partial equilibrium and is formulated in a life-cycle framework. Time is divided into discrete periods of one-year length. Agents enter the model at age 20 and die with probability one before turning 100. We assume that the agent works the first 45 years and retires afterwards.

2.1 Preferences and the labor income process

Households have Epstein-Zin utility function defined over one single non-durable consumption good. Letting U_t be household's utility at age t this can then be written as:

$$U_t = [c_t^\gamma + \beta(E_t U_{t+1}^\alpha)^\frac{\gamma}{\alpha}]^\frac{1}{\gamma} \quad (1)$$

In equation 1, β is the subjective discount factor, γ controls the elasticity of intertemporal substitution and α controls risk aversion. In particular the elasticity of intertemporal substitution is $\frac{1}{1-\gamma}$ and risk aversion is $1 - \alpha$. The expectation operator is taken with respect to uncertain labor earnings, stock

returns and survival. The latter is described by a time varying probability π_{t+1} of surviving up to age $t + 1$ conditional on being alive at age t .

The agent efficiency as a worker is age dependent according to the function $G(t)$. This function is meant to capture the hump-shaped profile of earnings over the working life. The deterministic component of labor efficiency units is hit by a stochastic shock represented by a first order autoregressive process in logarithms. Denoting the stochastic component of income with z_t this will then evolve according to the law of motion:

$$\ln z_t = \rho \ln z_{t-1} + \varepsilon_t \quad (2)$$

where ε_t is a normal i.i.d. shock. Following Cocco et al.(2005) we also allow for the possibility of a disastrous labor earnings shock, so that the labor income shock is near zero with a very small probability and follows equation 2 otherwise. We normalize wages to one so that labor income can be written simply as $y_t = G(t)z_t$. After retirement the agent receives a fixed pension benefit y_R^{zR} related to her earnings in the last working period, so that her nonfinancial income is $y_t = y_R^{zR}$.

2.2 Assets

Earnings shocks cannot be insured due to missing markets. The agent then uses savings to smooth consumption in the face of earnings fluctuations. In

doing so she has access to two assets. The first asset is a risk-free, liquid financial asset. This asset is meant to represent cash, checking and savings accounts, certificates of deposits and money market mutual funds, that is, all assets that are typically classified as liquid financial assets - as opposed to bonds and stocks - in the empirical literature. Wages are paid in the form of this asset which on top is the only asset that can be used to purchase consumption. We denote with m_t the amount of this asset that the agent holds at the beginning of period t and with R_{t+1}^m the return on holding the asset from time t to time $t + 1$. The second asset is a less liquid financial asset that we call stock for convenience. This asset is risky and provides a positive expected return premium above the liquid asset. This asset cannot be used directly to purchase consumption goods. We denote the amount of stock held at the beginning of period t with s_t and the return on holding stock from t to $t + 1$ with \tilde{R}_{t+1}^s . A no borrowing and no short-sale constraints are assumed.

The two assets are held in separate accounts and a fixed cost must be paid to make a transaction between the two accounts. This cost is fixed in the sense that it is independent of the amount of the risky asset that is traded. We make it proportional to earnings though, so as to capture the idea that the cost includes the monetary equivalent of the time spent to make financial

decisions.⁶ We denote the transaction cost with TC in the model. This is the key assumption in the model since it makes money more valuable as an asset to insure against consumption fluctuations.

Finally, we omit an explicit modeling of housing wealth given that this is not the focus of the model and would further complicate the numerical solution of the household's optimal program. However given the importance that housing has in households' economic decisions we decide to model it following the approach in Gomes and Michaelides (2005), who introduce in their model a flow of expenditures on housing services that does not give utility and that must be subtracted from income. We denote the fraction of income that is spent on housing with $h(t)$ to capture its dependence on the household's age.

2.3 Household's optimal program

Given the informal description of the individual problem stated above it is possible to write the household's optimization problem in dynamic programming form.

In describing the value function we first write the indirect utility in the case when the household decides to make a transaction between the two

⁶It is customary in the literature that uses entry costs to make them proportional to income; see for example Gomes and Michaelides (2005).

accounts. This will read:

$$V_t^{tr}(s_t, m_t, z_t) = \max_{c_t, s_{t+1}^o, m_{t+1}^o} \left\{ c_t^\gamma + \beta \mathbf{E}[V_{t+1}(s_{t+1}, m_{t+1}, z_{t+1})^\alpha]^\frac{\gamma}{\alpha} \right\}^\frac{1}{\gamma} \quad (3)$$

under the following constraints:

$$c_t + s_{t+1}^o + m_{t+1}^o \leq y_t(1 - h(t)) + m_t + s_t - y_t TC \quad (4)$$

$$s_{t+1} = \tilde{R}_{t+1}^s s_{t+1}^o, \quad m_{t+1} = R_{t+1}^m m_{t+1}^o \quad (5)$$

$$m_{t+1}^o \geq 0, \quad s_{t+1}^o \geq 0 \quad (6)$$

and the law of motion of z_t in equation 2. In this case the maximization of the right-hand side of the value function is taken with respect to consumption and both assets. Equation 4 is the budget constraint. The agent pays the fixed cost $y_t TC$ which allows her to buy or sell stocks, hence the amount of resources potentially available for consumption and asset purchases subtracts this cost from the sum of current earnings net of housing expenditures, money and stocks. The agent can then use these resources without further restrictions to buy consumption and the two assets. Equation 5 shows the laws of motion of stock and liquid holdings: It gives us the amount of resources in the monetary and stock accounts that the agent will have at the beginning of the next period, given the optimal choices of the two assets m_{t+1}^o and s_{t+1}^o . The last equation is the non-negativity constraint that applies to the holdings of the two assets. It simply says that the agent cannot short-sell

either asset. We use a separate notation for the control variables m_{t+1}^o and s_{t+1}^o and their corresponding state variables m_{t+1} and s_{t+1} because the return earned on the two assets makes the value of the control and state different.

Next we write the indirect utility in the case the agent decides not to perform any transaction between the money and stock account:

$$V_t^{ntr}(s_t, m_t, z_t) = \max_{c_t, m_{t+1}^o} \left\{ c_t^\gamma + \beta \mathbf{E}[V_{t+1}(s_{t+1}, m_{t+1}, z_{t+1})^\alpha]^\alpha \right\}^{\frac{1}{\gamma}} \quad (7)$$

subject to the following constraints:

$$c_t + m_{t+1}^o \leq y_t(1 - h(t)) + m_t \quad (8)$$

$$s_{t+1} = \tilde{R}_{t+1}^s s_t \quad (9)$$

$$m_{t+1} = R_{t+1}^m m_{t+1}^o \quad (10)$$

$$m_{t+1}^o \geq 0 \quad (11)$$

and equation 2. In the value function equation m_{t+1}^o denotes the amount of the liquid asset to carry into the next period. Equation 8 is the budget constraint. It reflects the fact that if no transaction between the two accounts is made the agent does not pay any fixed transaction cost but she will only be able to use her current earnings and the initial amount of money to purchase consumption. At the same time the balance on the monetary account carried over to the next period cannot exceed the sum of earnings net of housing expenditures and current money balances minus consumption. Equation 9

describes the fact that in the no transaction case the amount of stock carried to the next period is simply the gross return on the current amount. For this same reason in the equation defining the value function, maximization is taken only with respect to consumption and the liquid asset. Finally the last equation represents the usual no borrowing constraint.

As the laws of motion of stocks, equations (4) and (9) suggest, an implicit assumption is that either all the return on the stock takes the form of price appreciations or that dividends are immediately reinvested in the stock account at no cost. In reality part of the return on equity comes from dividends that are paid in the monetary account. Contrary to the standard model, with fixed transaction costs the way the return is split between capital gains and dividends is relevant for the investor's decision problem. For this reason in the result section we will also consider sensitivity analysis using an alternative version of the model where part of the return is paid in the form of a dividend.

The optimal value function and the optimal decision about whether to make a transaction or not is obtained by comparing the indirect utility in the two cases. This is summarized by the equation:

$$V_t(s_t, m_t, z_t) = \max\{V_t^{tr}(s_t, m_t, z_t), V_t^{ntr}(s_t, m_t, z_t)\}. \quad (12)$$

The model does not admit analytical solutions and is then solved numer-

ically. The solution to the model is especially difficult in this case for two reasons. First, once the fixed transaction cost is introduced the holdings of the two assets enter separately as a state variable, hence the model has three continuous state variables and two continuous controls.^{7,8} Second, the fixed transaction cost breaks the concavity of the objective function forcing the use of slow direct search methods for the optimization at each state space point.⁹ Details of the solution algorithm are provided in the Appendix.

3 Parameter Calibration

In the baseline simulation we set α to -4 and γ to -3, corresponding to a value for risk aversion and the elasticity of inter-temporal substitution equal to 5 and 0.25 respectively. The subjective discount factor β is set equal to 0.94, a value that falls within the estimates by Gourinchas and Parker (2002) . The deterministic component of labor earnings $G(t)$ is represented by a third order polynomial. The coefficients of the polynomial are taken from the profiles estimated by Cocco et al. (2005). As far as the idiosyncratic shock is concerned we assume that it can be represented by an AR(1) process

⁷Models without a fixed transaction cost only have two continuous state variables, that is, the sum of all financial assets and (lagged) income.

⁸Following the standard approach in solving these models the labor shock, while it is theoretically continuous, will later be approximated with a discrete Markov process.

⁹See Corbae (1993) on this point.

in logarithms, that is, we assume $\ln(z_{t+1}) = \rho \ln(z_t) + \varepsilon_{t+1}$ where ε is a normal random variable $N(0, \sigma_\varepsilon^2)$ and is i.i.d. We use the estimates of the process in Hubbard et al. (1995) who find values of ρ of 0.946 for high school graduates and 0.955 for college graduates and values of σ_ε^2 of 0.025 and 0.016 respectively.

Retirement income includes two components: a social security benefit and a pension benefit. In calibrating the social security component we apply the formula used in the U.S. system to compute old age retirement benefits, as described for example in Huggett and Ventura (2000). This formula is based on the computation of an average of monthly earnings during working-life, called AIME. The social security payment is then obtained by applying a replacement ratio of 90 percent up to 0.2 times average life-time earnings, a marginal replacement ratio of 32 percent from 0.2 to 1.24 times average life-time earnings and a marginal replacement ratio of 15 percent above 1.24 times average life-time earnings. No further benefit is credited above 2.47 times average life-time earnings. A strict application of this formula would require the addition of a further continuous state variable. In order to avoid that we exploit the high persistence of the earnings process and compute mean life-time earnings conditional on the shock in the last year of work. We then compute the population average of this measure of life-time earnings and we apply the formula used by the U.S. social security system to compute

the benefit. While not fully linking benefits to working age average earnings, our calibration still introduces some progressive features that help matching wealth-to-income ratios across the wealth distribution. Munnell and Soto (2005) report a median replacement ratio including defined benefit pensions of 60 percent for single and 67 percent for couples. We thus add a component that is proportional to earnings in the last year of working life so that the median replacement ratio is 61 percent. As for the housing expenditure process we assume that it is described by a third order polynomial and take the values of the coefficients from the estimates presented in Gomes and Michaelides (2005).

The real return on the liquid asset is 2 percent and that the expected real return on the stock is 6 percent. Following a tradition in this literature, the implied premium is lower than the historical one.¹⁰ The process for the stock return is assumed to be normal and i.i.d. over time with a standard deviation of 18 percent, in line with the historical evidence about the US Standard and Poor's 500 index. We calibrate the initial wealth distribution using data from the Survey of Consumer Finances. More specifically we match the first two moments of the empirical distribution for the youngest age group.

The most critical parameter to calibrate for the purpose of this model is the size of the transaction cost. This transaction cost includes both the

¹⁰See Cocco et al. (2005) for the reasons behind this choice.

monetary cost and non-monetary costs.¹¹ Lacking an adequate measure for the non monetary component of the cost we follow an alternative calibration strategy. Clearly the size of the cost will affect the frequency of transactions. We thus calibrate the cost so that once we simulate the model, the fraction of households that do not make a transaction in any given period matches the one in the data. To our knowledge there are two sources of data about households' transactions in the stock market. One are the reports "Equity Ownership in America" compiled by the Investment Company Institute and based on interviews of a sample of stock holding households. The second is the paper by Biliias et al. (2010) which reports data based on the PSID. The two sources report quite different figures. According to the report "Equity Ownership in America" about 40 percent of stockholding households make a transaction in a given year. According to Biliias et al. (2010) between 25 and 30 percent of the general population make a transaction over a 5 years period.

After the presentation of the simulated results we provide a detailed discussion of the calibration of the cost based on matching the empirical frequency of transactions within the model. This exercise will lead us to con-

¹¹Non-monetary costs include the time cost of gathering the information about the different assets and to make the decision about how much to invest in each, and "psychological" costs such as those required to overcome status quo biases or inertia.

sider two different levels of the fixed cost. In what follows these two different choices will be referred to as the low and high transaction cost case. In the low transaction cost scenario the calibrated value is 0.33 % of the household's non-financial income while, in the preferred high transaction cost scenario, the values are 4% for college graduates and 7 % for high school graduates. As we will see the main qualitative features of the model results that we want to highlight are common to the two levels of the cost, even though quantitatively the results will differ across the different experiments.

Finally we assume that the transaction cost is the same both for stock purchases and for stock sales. One might argue that the planning cost in the case of sales is much lower since an agent who needs to liquidate the asset in the face of negative earnings shocks or to supplement retirement income may simply do that with no planning. However, tax considerations may enter the sale decision and it might be natural to expect a high transaction cost in the case of gains but a lower transaction cost at least for losses because of the deferral of the tax liability — as studied for example in Gallmeyer et al. (2006). Empirically these effects appear to be small though. The data from both the Investment Company Institute (2002) and Biliias et al. (2010) suggests that sales occur slightly less frequently than purchases, both during market upswings and during market downswings.

4 Results

In this section we describe the results of the model. The section is divided into three subsections. In the first one sample decisions rules are reported. In the second subsection we report the results of the simulation of the model using one education group, namely the college educated. In this subsection we search for a baseline parametrization and check the robustness of the results to changes in preference parameters and the stock return process. In the third subsection we report the main results of this research. In this subsection we simulate the results for both education groups — college and high school educated — aggregate them and compare them to the empirical figures that refer to the aggregate population.

4.1 Decision rules

We report the optimal share invested in stock and the decision to make a transaction in the high transaction cost parametrization for an agent who is 45 years old. We start in figure 1 with a sample transaction decision. This decision rule refers to a high school agent with the lowest earnings shock.

[Locate figure 1 about here]

On the two horizontal axis we report the state variables, that is, current wealth and the share of this wealth invested in stocks prior to making the

decision.¹² On the vertical axis we report the decision to make a transaction. This decision is a discrete one and we make the convention that a 0 means that no transaction is made, a +1 that the agent buys stocks and a -1 that the agent sells stocks. The figure shows that for any level of wealth the agent will buy stocks when the current share is low, she will sell stocks when the current share is high and will not make any transaction in an intermediate range of the current share. The inaction region forms a band around the optimal share at each level of wealth. This band is large when wealth is small and narrows down when wealth increases since with more wealth a smaller deviation from the optimal share will make it convenient for the agent to pay the fixed transaction cost and readjust her portfolio.

[Locate figure 2 about here]

We next move to optimal share decision rules. We do that for a high school graduate with the lowest earnings shock and for a college graduate with the highest earnings shock. These two choices represent the two extreme values of human wealth endowment conditional on age. In figure 2 we examine the optimal stock share decision rule of the high school graduate with the lowest

¹²In the section describing the model the two state variables were the quantities of the two assets. The reasons for this change of variables are related to the numerical method used to solve the model and are highlighted in the appendix. Redefining the state variables is also more instructive for the purpose of understanding the mechanics of the model.

earnings shock. There are three main patterns that we want to highlight. In the transaction region the stock share is increasing in wealth at low wealth levels but once wealth passes a certain threshold the optimal share becomes decreasing with further increases in wealth. In the no transaction region, corresponding to the band in the middle of the graph, the optimal share is constant or mildly declining in wealth for a given current share and is increasing in the current share for a given wealth.

The interpretation of these patterns is the following. Because of persistence, a low earnings agent will want to hold some of her wealth in the form of the liquid asset in order to increase her consumption beyond her earnings without having to incur the fixed transaction cost. Given the amount of the liquid asset that is needed to accomplish this task, its share will decline with total wealth, hence the optimal stock share will increase. Past a certain level of wealth though, the optimal stock share will start to decline for the usual diversification reasons well highlighted for example in Cocco et al. (2005). In the no transaction region the forces at play are different and the optimal share is entirely determined by the total amount of stock at the beginning of the period and the optimal saving decision.

[Locate figure 3 about here]

In figure 3 we report the optimal decision rule for a 45 year old college graduate endowed with the highest earnings shock.¹³ In this case the graph can be divided in two broad areas. The first one corresponds to the no transaction region. In this region the optimal stock share is increasing in wealth for a given current share of stocks in the portfolio. As in the previous case this pattern arises from the interaction of the optimal saving decision and the current amount invested in stock. The second area corresponds to the region where the agent finds it optimal to pay the transaction cost. In this area the optimal stock share is equal to slightly less than 100 percent at low levels of current wealth and then declines. This pattern is similar to the one observed in standard models without transaction costs.

The patterns observed in figures 2 and 3 are representative of the patterns that, during working-life, we observe at other ages and earnings shocks. In general as the agent ages we observe a transition from the pattern represented in figure 2 to the one represented in figure 3. Conditional on a given educational group the switch from one to the other pattern occurs earlier in the life-cycle as we move towards higher earnings shocks. Similarly, conditional on a given earnings shock the transition occurs earlier in the life-cycle for the college graduates than for the high school graduates. Also the transition oc-

¹³We omit the corresponding graph for the optimal transaction decision since it does not add any new insight with respect to the graph for the lowest earnings shock agent.

curs in two ways. In the no-transaction region, as the agent ages, the decision rule conditional on wealth moves from being more declining than in figure 2 to the pattern shown in figure 3. In the transaction region the transition takes place in the form of a shrinking wealth range where the wealth-share relationship is increasing and an ever sharper hump, slowly turning to the pattern observed in figure 3. Finally, during retirement the decision rule for the optimal stock share follows a pattern that is similar to the one in figure 2 at all levels of the pension benefit but the range of wealth levels where the conditional stock share is increasing is in this case smaller.¹⁴

Summarizing, while in the standard models with no transaction costs the decision rules for the optimal stock share are monotonically declining in wealth, once fixed transaction costs are considered a more diverse picture emerges.¹⁵ In particular we can see that for low earnings agents the relationship between wealth and the optimal share of wealth invested in stocks is increasing in a range of low wealth levels, those presumably experienced by low earners. On the other hand, provided they are in the no transaction

¹⁴Results for the decision rules at different ages and income shocks as well as those for different parametrization are available upon request.

¹⁵This statement about the basic model with no transaction costs is true under the assumption that labor earnings are not correlated with the stock return. Under a sufficiently large positive correlation a different result would hold, however positive and high correlation is not supported empirically. See Cocco et al. (2005) and Haliassos and Michaelides (2003) on this point.

region, a similar relationship between wealth and the optimal stock share can be observed also for the decision rules of high earnings agents. Whether this is sufficient to generate a positive cross-sectional relationship between wealth and the stock share of market participants depends on the path of wealth accumulation through the different regions experienced by agents with different earnings history. This can be discovered by simulating the model.

4.2 Baseline simulation

We simulate a cohort of 2000 agents across their 80 period long life-cycle. Since the realized path of stock returns may affect the observed pattern of stock-holding we repeat the simulation 50 times to smooth out these fluctuations. The main focus of the results will be the behavior of the stock share conditional on participation by wealth and age. We omit the analogous results concerning participation rates since it is already known that fixed costs can generate the patterns observed in the data. We simulate a cohort of college graduates and compare the results to the corresponding values in the data for the same educational category. We first report a baseline parametrization of preferences, in particular risk-aversion and the intertemporal elasticity of substitution, and the stock return process. We then report the results of sensitivity analysis on this set of parameters. In a separate subsection we also present a discussion of the size of the transaction cost.

4.2.1 A baseline parametrization

In this subsection we describe results for the baseline set of parameters. We set α to -4 and γ to -3, corresponding to a value for risk aversion and the elasticity of inter-temporal substitution equal to 5 and 0.25 respectively. We simulate a cohort of college graduates. We report results for the high and low transaction cost scenario. The transaction cost is respectively 6.1 and 0.33 percent of income. In the low cost scenario the participation rate is 95.7 percent and the portfolio share of stocks conditional on participation is 84.5 percent. When the transaction cost is raised to match the level of inactivity reported in Biliias et al. (2010), the participation rates plunges to a value of 61.8 percent and the conditional stock share to 70.4 percent. This decline reflects the liquidity motive for holding the risk-free asset. When it is costly to make and carry out stock market investment decisions, households will want to hold a larger percentage of their wealth in the form of the risk-free, liquid asset to smooth their consumption in the face of time-varying and uncertain earnings. The participation rate for college graduates in the data is 69.8 percent and the conditional stock share is 62.8 percent, hence in the high cost scenario the participation rate and conditional stock share lie within about ten percentage points of its empirical counterpart.¹⁶

¹⁶The participation rate and conditional stock share are taken from the Survey of Consumer Finances, (2007).

We next move to the simulated conditional stock shares by wealth levels. This is done in table 1 which reports the average share of the financial portfolio held in stocks, conditional on participation, by wealth quartiles and separately for the top 5 percent wealthiest households. For comparison we also report the corresponding figures taken from the 2007 Survey of Consumer Finances. These figures are computed only on the subset of the college graduates whose earnings process was used to generate the simulated data. As it can be seen the model generates a relationship that is positive at low to intermediate levels of wealth independently of the size of the cost. In the low cost scenario the conditional share moves from 66.3 percent for the bottom wealth quartile to 95.9 percent for the third quartile and then declines to 69.3 percent for the top 5 percent of the wealth distribution. The model thus cannot reproduce a monotonically increasing profile, although it can explain why the poorest households hold a smaller share of stocks than those in the next richer quartiles of the wealth distribution. This result is quite important since it has been particularly difficult to explain this fact so far. The main explanation in fact relied on a strong and positive correlation between earnings shocks and stock market return which has little empirical support.¹⁷ In the high cost scenario results further improve. The conditional

¹⁷Wachter and Yogo (2010) propose an alternative theory based on non-homothetic preferences. That theory is able to generate shares of risky assets that are increasing in wealth

share is modestly increasing over the whole range of quartiles, moving from 60.0 to 74.2 percent from the bottom to the top one. It then modestly declines to 69 percent for the top 5 percent of the wealth distribution. In the data the share of stock for market participants increases from 56.2 to 67.6 percent over the four quartiles of the wealth distribution and then slightly declines to 63.3 percent in the top 5 percentiles.

[Locate table 1 about here]

In the data, when we condition on age, the relationship that exists between net worth and the share of financial wealth invested in stocks weakens. In table 2 we thus report the share of wealth invested in stocks by stockholders conditional on wealth by ten year age groups. The table is organized in three panels.¹⁸ The top one reports data from the Survey of Consumer Finances.

[Locate table 2 about here]

The other two panels report the simulated results of the model with transaction costs in the low and high cost scenarios. As it can be seen results are within most age groups. However they do not report the relationship between wealth and the stock share for the whole population, that is, without conditioning on age.

¹⁸The model simulates the life-cycle over 80 periods meant to represent age 20 to age 99. In the table we do not report the statistics for the two oldest age groups to economize on space. The patterns of stock holding by wealth observed within these two age groups do not differ from those for the other groups.

broadly similar to those that do not condition on age. In the low transaction cost scenario in the second panel we see that in the first two age groups, that is, the one from age 20 to 30 and from age 30 to 40 the relationship is broadly increasing from the bottom to the top of the distribution. For the 40 to 50 and 70 to 80 age groups the relationship has again the inverted U shape that can be found for the general population. Finally in the two age groups before and around retirement it is monotonically declining. Results for the high transaction cost scenario are reported in the last panel of the table. Under this scenario, the share of wealth held in stocks is broadly increasing in wealth up to the 40 to 50 age group. For the remaining age groups they show an inverted U-shaped pattern but the declining leg is modest and allows the share of the top of the distribution to be similar to or higher than the share at the bottom. Once again this represents an improvement over the standard model where at all ages the relationship between wealth and the stock share is negative unless positive contemporaneous correlation between earnings shocks and stock returns is assumed, something not supported by the empirical evidence. Notice that even more recent models that exploit some more sophisticated form of correlation between earnings and stock market performance like the one of Lynch and Tan (2011) still would run into trouble for agents close to or past the retirement age when there is no or very little wage uncertainty remaining, hence little or no room for any

pattern of correlation between nonfinancial income and the stock return.

Finally in figure 4 we report the allocation to stocks along the life-cycle for stock market participants. The continuous line shows the empirical profile which exhibits an hump-shaped pattern. The dashed and dashed dotted lines represent the life-cycle profiles for the models with fixed transaction costs in the high and low cost scenarios. In both cases the pattern of stock shares is increasing in age in the first part of the working life. The share then declines to give rise to a hump-shaped trajectory in the low cost scenario, while it remains roughly constant in the high cost scenario. Overall the life-cycle profile for the high cost scenario follows quite closely its empirical counterpart, while the one in the low cost scenario is somewhat higher, a feature that could be already foretold from the life-cycle averages previously reported.

[Locate figure 4 about here]

One caveat is in order concerning these profiles. The empirical one is obtained as the cross-section of the observed stock share for stockholders. Estimation work conducted by Ameriks and Zeldes (2004) has shown though that the actual profile depends on the underlying identifying assumptions.¹⁹

¹⁹The issue arises because age, time and cohort are linearly dependent so that when constructing age profiles it is impossible to simultaneously identify time and cohort effects.

The profile can be either increasing or mildly hump-shaped depending on the identifying assumption. In light of this observation it is clear that while models without transaction costs gives rise to counterfactual results, since the higher share invested by young households is not observed in the data regardless of the controls in the estimation, the model with transaction costs can match the mildly increasing profile during working age that is observed in the data, again independently of the identifying assumptions.

4.2.2 Frequency of transactions and calibration of the transaction cost

In this section we discuss the size of the transaction cost and compare it with empirical evidence and with other calibrated models in the literature. In the low cost scenario, in order to obtain the frequency of transactions reported in the ICI survey we set the fixed cost at a level of 0.33 percent of the household's non-financial income. The simulated data show that in this scenario households perform on average one transaction every 1.95 years, so that the transaction cost paid per year of participation in the stock market turns out to be 0.17 percent of average income. In the high cost scenario, in order to match the frequency of transactions observed in the PSID and reported in Biliias et al. (2010) we need to set the cost to 6.1 percent of the household's non financial income. However, in this case the frequency of transactions in

the simulated data drops to one every 6.3 years of participation in the stock market. Consequently the transaction cost per year of participation is in fact only 1.27 percent of average income. For comparison we may look at the empirical estimates presented in Vissing-Jørgensen (2002). Those estimates refer indeed to participation costs, however we believe the comparison is still useful to get a sense of whether our values are of the correct order of magnitude. Using the 1989 and 1994 waves of the PSID Vissing-Jørgensen (2002) estimates that a per period participation cost of 260 dollars (in 2000 prices) is needed to rationalize the non-participation of three quarters of the households. Compared to average earnings of about 40000 dollars this is about 0.6 percent.²⁰ As we can see the per period amount of the transaction cost paid in the low cost scenario is below the figure reported by Vissing-Jørgensen. In the high cost scenario it is somewhat larger but still in the same order of magnitude.

In the macro literature, Khan and Thomas (2011) consider a model with two assets, money and bonds, to study the impact of monetary shocks. They calibrate the cost of making a transaction between the two accounts to match aggregate money velocity. In order to do that they assume a uniform distribution of the cost with support between 0 and 25 percent of the average endowment. Our values, as a percentage of income, fall well within that

²⁰The source for this figure is Díaz-Giménez et al. (1997)

range.

In addition Gomes et al. (2009) show that, for most households, their tax-deferred accounts are the main savings vehicle. But those tax-deferrable vehicles are largely illiquid, with significant penalties applicable in case of pre-retirement withdrawals. Moreover they show that, for about 40 percent of the population — which they call “indirect stockholders ”—, all equity is held exclusively inside these tax-deferred accounts. The liquid/transaction account has smaller balances and these are fully invested in riskless securities. Even the “direct stockholders ”, which have equity in both, invest a higher fraction of it within their retirement accounts, while keeping a buffer stock of wealth in a liquid account which has a more moderate equity allocation. Furthermore they also show that this can be partially motivated by the desire to smooth income shocks, which creates an incentive to keep assets with lower risk in the liquid account. As a result equities — both in the model and in the data — become subject to a higher effective rebalancing cost.²¹ This provides an additional rationale for the calibration of the transaction costs associated with the equity account, within our model. In addition, the “indirect stockholders ”— for whom this constraint is more binding — have

²¹Ameriks and Zeldes (2004) show that, in a large sample of households with tax-deferred accounts, 44 percent of those households made no changes whatsoever to either their flow or asset allocations over the ten-year period, and another 17.2 percent made only one change to either stocks or flows.

significantly less total wealth than the “direct stockholders ” and as a result they are the more relevant ones at low levels of wealth, which is exactly the region where the decision rule for the portfolio share in the model with transaction costs is particularly steep.

Overall we believe it is reassuring that our model matches the observed frequency of transaction with a cost that is comparable both with the values found in the empirical literature and with those used in quantitative models that follow similar calibration procedures.

4.2.3 Changing risk aversion

In this subsection we check the behavior of the model when risk aversion is reduced studying the effect both on the frequency of transaction and the behavior of stockholding over the life-cycle and the wealth distribution. We perform this exercise both for the high and for the low cost scenario.

We reduce the level of risk-aversion while keeping the elasticity of substitution constant. This change in preference parameters leads to a reduction in wealth accumulation. Since the size of the inaction region is greater at lower wealth levels this change in parameters has the potential to reduce the frequency of transactions. We then need to change the size of the transaction cost to still match the empirical target. The parameter α is increased to -1 corresponding to a coefficient of risk aversion of 2.

In the high transaction cost case the results confirm the intuition laid down above and indeed we need to reduce the size of the transaction cost from 6.1 to about a third the original value, that is, 2 percent of income. In the low transaction cost scenario, perhaps surprisingly we indeed need to increase the size of the transaction cost, from 0.33 to 0.39 percent of income in order to match the target of 40 percent of stockholders making a transaction in a year. In order to understand this result we conjecture that a lower risk aversion increases the benefits of holding stocks, everything else equal reducing the size of the no transaction region.

At the aggregate level the participation rate in the high cost scenario increases to 62.6 moving a bit closer to the 69.8 that we observe in the data. The conditional share increases from 70.4 percent in the baseline simulation to 74.9 percent, further away from the 62.8 percent in the data. In the low cost scenario the increase in the transaction cost, coupled with the decrease in wealth accumulation reduces the participation rate to 89.5 percent. The conditional stock share in this case increases to 86.7 percent. The increase in the stock share in both cases is the consequence primarily of the reduced risk aversion. To a lesser extent it is also a consequence of the reduced wealth accumulation which reduces the frequency with which households move through the regions where the optimal stock share becomes declining in wealth.

We report stock shares by wealth quartiles and by age for stockholders in table 3 and in figure 5 where for comparison we report again the figures for the college graduate group in the data.

[Locate table 3 about here]

As we can see in the low transaction cost case the conditional share becomes monotonically increasing and ranges from 56.9 percent in the bottom quartile to 98.8 percent at the top 5 percent. In the high cost scenario we observe a decrease in the share from 69.1 to 55.1 percent between the bottom and next quartile, followed by a monotonically increasing pattern that tops at 92.5 percent.

The life-cycle profiles of the stock share are reported in figure 5 together with the analogous data for the college graduates in the 2007 Survey of Consumer Finances. As we can see in the low cost scenario the life-cycle pattern of stock shares is clearly hump-shaped like in the data, even though at a quite higher level. In the high cost scenario the profile shows a reduction in the stock share for market participants between the first and second age group. Apart from that it is hump shaped like in the data and quantitatively close to it.

[Locate figure 5 about here]

Overall these experiments suggest that a suitable reduction in the coefficient of relative risk aversion allows the model to match the fraction of transacting agents with a substantially lower cost. The change in the cost is not an order of magnitude change though. Moreover the effect is present only when the cost is high to start with. Overall the other properties of the model are preserved under this parameter change but the conditional share moves slightly away from the target data.

[Locate table 4 about here]

4.2.4 Positive correlation between earnings and stock return

In this section we present results for a version of the model where we assume that the shock to earnings shows positive contemporaneous correlation with the stock return. Since results do not change much compared to the baseline and to economize on space in this case we only consider the baseline preference parameters and the high transaction cost scenario. We select the high cost scenario because it generates a participation rate that is close to the one in the data, while the low cost scenario generates participation rates that are more than 20 percentage points above the empirical counterpart. Moreover the average portfolio allocation to stocks in the population of stockholders is also closer to the data making the comparison with the latter more significant.

We set the value of the correlation to 0.15, the number estimated and used by Campbell et al. (2001). As expected the fraction of the portfolio allocated to stock in the population of stock holders goes down, falling from 70.4 percent to 65.4 percent. The behavior of the stock shares by wealth levels is reported in table 4. Table 4 shows that under the assumption of positive correlation between stock returns and labor earnings shocks, the pattern of stock shares shows the usual hump-shaped pattern, with an increase in the share from 59.3 percent to slightly more than 70 percent between the bottom and the top quartile followed by a decline to 65.4 percent for the top 5 percent of the distribution. The dashed dot line in figure 6 reports the life-cycle profile of the conditional stock share in this case. This profile is slightly increasing from an average share of about 60 percent in the first decade of working life to about 70 percent before retirement and after. It almost coincides with the empirical one — represented in the graph by the continuous line — in the first two decades but then exceeds it by a small amount in the decades around and after retirement.

[Locate figure 6 about here]

4.2.5 Separate dividend and capital gain

In this section we consider a stock return process in which part of the stock return is paid as cash dividend. We simulate the model for the baseline set

of preference parameters and the high cost scenario.

Most life-cycle asset allocation models do not make a distinction between the capital gain and the dividend component of the return on public equity and for this reason we decided to follow this tradition in the baseline model. Absent fixed transaction costs this practice is irrelevant for the results. In the presence of fixed transaction costs though this is not any more true. If stocks pay a dividend on the liquid account they become themselves a source of money that can relax the cash-in-advance constraint. This would make stocks more attractive compared to the case where all the return is in the form of price appreciation. On the other hand to the extent that the transaction cost makes transactions from the liquid to the stock account infrequent, the dividend might sit for several periods in the liquid account yielding a lower return than the one it would earn on average in the stock market. This in turn would make stock less attractive. To check the implications of a separate dividend and capital gain component of the return we now assume that stock holdings pay a constant 2 percent dividend and that price growth is stochastic and averages 4 percent. The standard deviation of the price appreciation is set at 18 percent. This calibration is meant to leave both the expected return and its volatility unchanged from the baseline model. This choice overestimates the contribution of the dividend to the overall stock return. For example, Dammon et al. (2004) use a 2 percent nominal

dividend yield in a model where the capital gain is set at 9 percent. The average participation rate is in this case 61.5 percent and the average share of the portfolio allocated to stocks by stock market participants is 69.3 percent, both values are very close to the ones in the baseline model. Results for the pattern of stock shares by wealth are reported in table 4. The pattern is now a bit more markedly hump-shaped: the conditional share rises from 58.9 for the bottom quartile to 71.8 percent for the top quartile of the wealth distribution but then declines to 63.6 percent in the top 5 percentiles. The dashed line in figure 6 reports the conditional stock share by age. As in the previous case the pattern is mildly increasing from a bit below 60 percent in the first two decades of working life to a bit more than 70 percent starting in the 50 to 60 age group. Summarizing we can say that the two potential effects of separating the dividend yield from the capital gain component of the return offset each other leaving investors' behavior almost unchanged even if the assumed dividend yield somewhat overstates the empirical one.²²

²²We also simulated the model under the assumption of a 1 percent dividend yield. This assumption would make the contribution of the dividend yield to the total return on stocks equal to the one in Dammon et al. (2004). Results were even closer to the ones of the baseline parametrization. For this reason we do not report them in the paper.

4.3 Analysis of the full model

In the previous discussion we focused on the behavior of a population made by households facing the earnings process of college graduates. In the current subsection we will consider a more general population composed of two educational groups, college and high school educated households in the proportion of 37 and 63 percent.²³ We report the usual statistics about the stock share by wealth and age cells and also, in the next section, some statistics relative to wealth-to-income ratios by age groups and quartiles of the wealth distribution. We view this as the main result of the paper since the aggregation of heterogeneous education groups allows us direct comparison with the aggregate data on the conditional stock share by wealth and over the life-cycle that have been the object of much of the literature.

4.3.1 Calibration of the transaction cost

We work with the high transaction cost case. However we notice that keeping the transaction cost constant at the level used in the benchmark simulation leads to a participation rate that is below the empirical value for the college graduates and above the empirical value for the high school group. In particular under this scenario the model would generate a participation rate that is

²³The proportion of the two groups are calculated based on the 2007 Survey of Consumer Finances.

65.6 percent for college graduates and 62.4 percent for high school graduates to be compared with 69.7 and 39.8 percent in the data. It would also generate a conditional stock share that is 70.9 percent for college graduates and 69.9 percent for high school graduates. In this dimension then, the difference in the two groups is smaller than in the data since in the data the share is 62.8 percent for college graduates and 59 percent for high school graduates. Also the percentage of agents that make a transaction over a five year period is virtually identical at the target value for both educational groups. This is because the cost is proportional to income and the same factor of proportionality is applied to the two groups. The only possible difference could be a different timing of wealth accumulation because of the different shape of the earnings profile but as it turns out the effects are negligible. Overall these results confirm the difficulty of obtaining substantial differences across educational groups in this kind of models already noticed in Cocco et al. (2005).

For this reason we allow the cost to differ across groups and calibrate it to target the average fraction of agents in the population that does not make a transaction over five years from Biliias et al. (2010) and the stock market participation rate for college graduates that we compute from the 2007 Survey of Consumer Finances. The resulting transaction cost is 4 percent of annual income for college graduates and 7 percent of annual income for the high

school graduates.

[Locate table 5 about here]

4.3.2 Simulation results

The model generates in this case a participation rate of 66.1 percent and a portfolio share of stock for stockholders of 70.0 percent. The corresponding figures for the whole population in the 2007 Survey of Consumer Finances are 51.1 percent for participation and 60.1 for the conditional stock share. The conditional stock share is then still within 10 percentage points of the empirical value while the participation rate exceeds the empirical one by more than 15 percent mainly because, despite the substantial transaction cost, participation for the high school graduates is quite higher in the model than in the data. Indeed in the model the participation rate for the college graduates is at the target of 69.8 percent, while for the high school graduates it declines to only 63.1 percent versus the 39.8 percent in the data reported above. With respect to the conditional stock share in the model it is 68.1 percent for the high school graduates and 73.2 for the college graduates. The difference is now 5 percentage points, close to the 3.8 percentage points difference in the data. Finally the fraction of agents that make a transaction in any given 5 years period now increases for the college graduates to 31 percent while it declines slightly for the high school graduates. This fact

is consistent with regression analysis reported in Biliias et al. (2010) that shows that the probability of being active in the stock market increases with education.

Conditional shares by wealth percentiles are reported in table 5 for the general population both in the model and in the 2007 Survey of Consumer Finances.²⁴ In the latter the profile is mildly increasing with a portfolio share of stock that increases from 55.9 percent in the bottom quartile to 62.5 percent in the top 5 percent of the distribution. A similar pattern is observed in the model but the increase is larger: from 57.8 percent in the bottom quartile to 72.8 percent in the top five percent of the distribution. Moreover there is a very mild hump in the profile since the share invested in stocks tops at the top quartile with a value of 75.5 percent.

Finally the life-cycle profiles of conditional stock shares are reported in figure 7 where the continuous line represents the data and the dashed line represents the model generated profiles.

[Locate figure 7 about here]

The data show a hump shaped pattern. The profile in the model follows very closely the one in the data for the first three decades of life but then it

²⁴To economize on space we do not report these figures separately for the two groups. However they are qualitatively very similar to each other and to the patterns of the whole population.

keeps increasing albeit at a slow pace and flattens out past the age 60 at a value of about 74 percent.

4.3.3 Wealth accumulation and wealth distribution

In this section we report the wealth-to-nonfinancial income ratios at the percentiles of the wealth distribution and by age groups that correspond to those that we have used to describe the patterns of conditional stock shares. Since an important focus of the present paper is to show that the introduction of fixed transaction costs improves the ability of the model to explain the relationship between stock shares and wealth, it is important that the wealth accumulation pattern generated by the model is broadly consistent with the one that is found in actual data.

We perform this exercise using the model that simulates the population of college and high school educated investors jointly and comparing it to the corresponding empirical data for the whole population.

[Locate table 6 about here]

Results are reported in table 6. The empirical figures are constructed using financial wealth as a measure of the household's wealth. Since in the present paper there is no explicit housing wealth and housing expenditures reduce the amount of income available for consumption and savings, this is

the more appropriate data counterpart to the model.

As we can see in the top panel of table 6 the model provides a good approximation to the data, along the wealth distribution. The wealth-income ratio generated by the model is slightly above the one in the data at the 25th percentile of the distribution and slightly below at the 95th percentile. The model still provides a reasonable approximation at the other two percentiles although both at the 75th percentile and especially at the median it somewhat overestimates the wealth-income ratio that we find in the data.

In the bottom panel of table 6 we report the wealth to income ratio by 10 years age groups. The first row shows the empirical figures and the second row the corresponding figures generated by the model. The model produces values that are very close to their empirical counterparts early in life. They are somewhat higher during mid-life especially in the retirement decade. However, in this class of models, given the absence of certain additional reasons for saving late in life, like medical or long term care risk, wealth to income ratios will fall somewhat more rapidly during retirement, than what we observe in the data.²⁵ Therefore, to avoid having large differences in wealth accumulation during that period, we overshoot to some extent the wealth to income ratios in the age brackets prior to retirement.

²⁵See Dynan et al. (2004) about the role that precautionary savings against medical expenditures plays in boosting wealth holdings, after retirement.

As we can see the wealth to income ratio is then only slightly below the data for households aged between 70 and 80 years.

5 Summary and Conclusions

In the current paper we have constructed a life-cycle portfolio choice model with uninsurable labor earnings risk. There is by now an important literature in this area. The current paper departs from that literature in that it re-interprets the risk-free asset as a liquid financial asset. Consumption can be purchased only with the liquid asset and a cash-in-advance constraint that can be relaxed by paying a fixed-cost to make transactions between the stock and the liquid account is assumed.

The assumptions made allow it to produce some improvements over conventional models that assume entry or participation costs. These improvements are in the area of the allocation to stocks over age and wealth for households who participate in equity markets. These results are obtained using parameterizations that are also consistent with empirical population averages. In particular the model generates average participation rates and conditional stock shares that are only a few percentage points above their data counterpart. The fact that the model slightly over-predicts the stock share is not surprising. The model abstracts from certain sources of back-

ground risk like marital and health risk that would help reduce the share allocated to stocks.²⁶ Integrating these features into the model represents an interesting and promising avenue for future research.

References

- [1] Abel, A.B., Eberly, J.C., Panageas, S., 2007. Optimal Inattention to the Stock Market. *American Economic Review*, 97, 244-249.
- [2] Akyol, A., 2004. Optimal Monetary Policy in an Economy with Incomplete Markets and Idiosyncratic Risk. *Journal of Monetary Economics* 51, 1245-1269.
- [3] Alvarez, F., Guiso, L., Lippi, F., 2012. Durable Consumption and Asset Management with Transaction and Observation Costs. *American Economic Review* 102, 2272-2300.
- [4] Alvarez, F., Atkeson, A., and Kehoe, P.J., 2002. Money, Interest Rates and Exchange Rates with Endogenously Segmented Markets. *Journal of Political Economy* 110, 73-112.

²⁶The impact of demographic shocks on asset allocation decisions has been studied by Love (2009).

- [5] Ameriks, J., Zeldes, S., 2004. How Do Household Portfolio Shares Vary with Age. Mimeo.
- [6] Ang, A., Papanikolaou, D., Westerfield M.M., 2011. Portfolio Choice with Illiquid Assets. Mimeo.
- [7] Benzoni, L., Collin-Dufresne, P., Goldstein, R.S., 2007. Portfolio Choice over the Life-Cycle when the Stock and the Labor Market are Cointegrated. *Journal of Finance* 62, 2123-2167.
- [8] Bertaut, C.C., Starr-McCluer, M., 2000. Household Portfolios in the United States. Federal Reserve Board of Governors, Finance and Economics Discussion Paper, 26.
- [9] Biliass, Y., Georgarakos, D., Haliassos, M., 2010. Portfolio Inertia and Stock Market Fluctuations. *Journal of Money, Credit and Banking* 42, 715-742.
- [10] Bonaparte, Y., Cooper, R., Zhu, G., 2012. Consumption smoothing and portfolio rebalancing: The effects of adjustment costs. *Journal of Monetary Economics* 59, 751-768.
- [11] Brunnermeier, M.K., Nagel, S., 2008. Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation. *American Economic Review* 98, 713-736.

- [12] Calvet, L.E., Campbell, J.Y., Sodini, P., 2007. Down or Out: Assessing the Welfare Costs of Household Investment Mistakes. *Journal of Political Economy* 115, 707-747.
- [13] Campbell, J.Y., Cocco, J., Gomes, F. and Maenhout, P., 2001. Investing Retirement Wealth ? A Life-cycle Model. In: Feldstein M., Campbell J.Y. (Eds.), *Risk Aspects of Investment-Based Social Security Reform*. University of Chicago Press. Chicago, IL.
- [14] Chiappori, P., and Paiella, M., 2011. Relative Risk Aversion Is Constant: Evidence from Panel data. *Journal of the European Economic Association* 9, 1021-1052.
- [15] Chien, Y., Cole, H., Lustig, H., 2012. Is the Volatility of the Market Price of Risk Due to Intermittent Portfolio Rebalancing? *American Economic Review* 102, 2859-2896.
- [16] Cocco, J., 2005. Portfolio Choice in the Presence of Housing. *Review of Financial Studies* 18, 535-567.
- [17] Cocco, J., Gomes, F., Maenhout, P., 2005. Portfolio Choice over the Life-cycle. *Review of Financial Studies* 18, 491-533.

- [18] Corbae, D., 1993. Relaxing the Cash-in-advance Constraint at a Fixed Cost. Are Simple Trigger-target Portfolio Rules Optimal? *Journal of Economic Dynamics and Control* 17, 51-64.
- [19] Dammon, R.M., Spatt, C.S., Zhang, H. H., 2004. Optimal Asset Location and Allocation with Taxable and Tax Deferred Investing. *Journal of Finance* 59, 999-1037.
- [20] Díaz-Giménez, J., Quadrini V., Ríos-Rull, J.V., 1997. Dimensions of Inequality: Facts on the U.S. Distribution of Earnings, Income and Wealth. *Federal Reserve Bank of Minneapolis Quarterly Review* 21, 3-21.
- [21] Dynan, K.E., Skinner, J. Zeldes, S.P., 2004. Do the rich save more? *Journal of Political Economy* 112, 399-444.
- [22] Gallmeyer, M., Kaniel, R., Tompaidis, S., 2006. Tax management strategies with multiple risky assets. *Journal of Financial Economics* 80, 243-291.
- [23] Gomes, F., Michaelides, A., 2003. Portfolio Choice with Internal Habit Formation: A Life-Cycle Model with Uninsurable Labor Income Risk. *Review of Economic Dynamics* 6, 729-766.

- [24] Gomes, F., Michaelides, A. 2005. Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence. *The Journal of Finance* 60, 869-904.
- [25] Gomes, F., Michaelides, A., Polkovnichenko, V., 2009. Optimal savings with taxable and tax-deferred accounts. *Review of Economic Dynamics* 12, 718-735.
- [26] Gourinchas, P.O., Parker, J., 2002. Consumption over the life-cycle. *Econometrica* 70, 47-89.
- [27] Guiso, L., Haliassos, M., Jappelli, T. (Eds.), 2001. *Household Portfolios*. MIT Press, Cambridge and London.
- [28] Haliassos, M., Michaelides, A., 2003. Portfolio Choice and Liquidity Constraints. *International Economic Review* 44, 143-177.
- [29] Heaton, J., Lucas, D., 1997. Market Frictions, Savings Behavior, and Portfolio Choice. *Macroeconomic Dynamics* 1, 76-101.
- [30] Heaton, J., Lucas, D., 2000. Portfolio Choice in The Presence Of Background Risk. *Economic Journal* 110, 1-26.
- [31] Huggett, M., Ventura, G., 2000. Understanding Why High Income Households Save More than Low Income Households. *Journal of Monetary Economics* 45, 361-397.

- [32] Investment Company Institute. 2002 and 2005. "Equity Ownership in America". (New York, NY and Washington, DC)
- [33] Khan, A., Thomas, J.K., 2011. Inflation and Interest Rates with Endogenous Market Segmentation. Mimeo.
- [34] Love, D.A., 2009. The Effects of Marital Status and Children on Savings and Portfolio Choice. *Review of Financial Studies* 23, 385-432.
- [35] Lynch, A.W., Tan, S., 2011. Labor Income Dynamics at Business Cycle Frequencies: Implications for Portfolio Choice. *Journal of Financial Economics* 101 333-359.
- [36] Munnell, A.H., Soto, M., 2005. What replacement rates do households actually experience in retirement? Mimeo.
- [37] Polkovnichenko, V. 2007. Life-Cycle Portfolio Choice with Additive Habit Formation Preferences and Uninsurable Income Risk. *The Review of Financial Studies* 20, 83-124.
- [38] Vissing-Jørgensen, A., 2002. Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures. NBER WP # 8884.
- [39] Wachter, J.A., Yogo, M., 2010. Why Do Household Portfolio Shares Rise in Wealth. *Review of Financial Studies* 23, 3929-3965.

- [40] Yao, R., Zhang, H., 2005. Optimal Consumption and Portfolio Choice with Risky Housing and Borrowing Constraints. *Review of Financial Studies* 18, 198-239.

Appendix A. Numerical solution.

In this Appendix we describe the numerical solution method. Since the simulation is standard we will focus our attention on the dynamic programming problem where the fixed transaction cost introduces certain complications that make the numerical solution harder and more time consuming. The household's maximization problem can be described by a finite horizon dynamic program that can be solved by the well known backward iteration method. The assumption of a fixed transaction cost in the stock account introduces two complications that must be addressed. The first one is the need to keep track separately of the amount of the two assets held. The second one is that a fixed transaction cost makes the value function non concave thus making fast optimization algorithms — like Newton methods — unsuitable. This latter point was made in a two period model by Corbae (1993).

To solve the dynamic programming problem we first make a variable transformation. To introduce this transformation, let first denote with \bar{s} and \bar{m} the upper bound of the stock and bond interval where the numerical value function is defined. With no borrowing and no short sale constraint the value function will thus be defined over the set $[0, \bar{s}] \times [0, \bar{m}]$. Let $M = \{m_1, m_2, \dots, m_n\}$ and $S = \{s_1, s_2, \dots, s_m\}$ be the grid points for the liquid and illiquid asset respectively with $m_n = \bar{m}$ and $s_m = \bar{s}$. At each iteration on

the value function the problem defined by the RHS of the Bellman equation must be solved for all pairs (s_i, m_j) with $s_i \in S$ and $m_j \in M$. Clearly for $s_i = \bar{s}$ and $m_j = \bar{m}$ or pair of sufficiently high values for the two state variables the constraint set includes choices for m^o and s^o that far exceed \bar{m} and \bar{s} imposing the evaluation of the value function by extrapolation far from the rectangle where it is defined. This may introduce severe approximation errors and for this reason we decided to redefine the value function over an alternative but equivalent set of state variables that include the current wealth W and the current share invested in stock that we denote with α . Under no borrowing and no short sale constraints α is bounded between 0 and 1 and no extrapolation of the value function is ever needed. Clearly extrapolation may be needed along the wealth dimension but in a way that is no different than in the standard consumption-saving model and hence to a much more limited degree than without the variable transformation.

To address the second problem, that is, the non concavity of the value function, we decided to use a two step direct search method. In the first step we define an action grid that is denser than the state space grid and search for the maximum over the action grid. That is, if $\alpha \in [0, 1]$ takes values $\{\alpha_1, \alpha_2, \dots, \alpha_n\}$ in the state space grid, for each interval $[\alpha_i, \alpha_{i+1}]$ we lay $n_\alpha - 1$ equally spaced points. Similarly, given $W \in [0, \bar{W}]$ for each interval defined by two adjacent points in the state space for wealth $[W_i, W_{i+1}]$ we lay

$m_w - 1$ equally spaced points. The maximization of the RHS of the Bellman equation then is performed by directly searching over the whole set of $((n - 1) * n_\alpha + 1) \times ((m - 1) * m_w + 1)$ points defined by the finer grid and not just on the $n * m$ points of the state space grid. This gives a first approximation to the solution, say the pair (α_{i^*}, w_{j^*}) . Next we refine the solution along the wealth dimension. We fix α_{i^*} , consider the two-sided interval around w_{j^*} , that is the interval $[w_{j^*-1}, w_{j^*+1}]$ and lay n_* points between the two extremes.²⁷ We then search over this new grid for the maximum. If we call this new maximum w_{j^*,j^*} , the solution to the maximization problem will be the pair $(\alpha_{i^*}, w_{j^*,j^*})$.

The state space grid contains 251 points along the wealth dimension and 41 along the current stock share dimensions. Points along the latter dimension are equally spaced, while those along the former are set so that the grid is finer close to the origin and becomes progressively coarser. Within each interval determined by the state space points we set 4 points along the conditional share dimension and 2 along the wealth dimension, that is, each interval is divided into 5 and 3 subintervals respectively. In the refined search along the wealth dimension we use 200 points in each $[w_{j^*-1}, w_{j^*+1}]$ subinterval. In order to evaluate the value function at points in the choice space

²⁷Clearly when w_{j^*} falls at the edge of the domain of the numerical value function along the wealth dimension the interval around w_{j^*} will be one-sided.

that do not coincide with points in the state space we interpolate by using bi-cubic spline approximation of the value function. The chosen grid search optimization methods makes the problem effectively discrete. Using a choice space that is finer than the state space and the two step search allow us to reduce the number of function evaluations while ensuring more accuracy in the solutions. To give an idea of the accuracy of the method, observe that the grid implies that the optimal choice of α is done over steps of 0.5 percentage points. The grid over wealth is non-uniform, hence a single number cannot be given. Around average earnings the step corresponds in economic terms to 0.02 percent of that average.²⁸ We tried to double the number of points along both dimension and did not find that changed the results in any significant way.

²⁸In dollar terms, if we assume an average wage of between 40000\$ and 50000\$ this corresponds to between 15 and 20 \$.

Appendix B. Data Construction.

In this appendix we briefly describe the procedure used to construct the empirical stock shares. Data come from the 2007 issue of the Survey of Consumer Finance. The Survey of Consumer Finance is a survey conducted every three years by the Board of Governors of the Federal Reserve System. It is widely used as a source for data about households' balance sheet since its non-random design offers more reliable information about the asset holdings at the top of the wealth distribution.

In order to construct portfolio stock shares for stock market participants by wealth we classify households based on quartiles of net worth. Net worth is defined as the sum of financial and non-financial assets minus all debt. Financial assets include all liquid accounts, certificates of deposits, stocks and bonds held both directly and indirectly, retirement accounts, the cash value of life-insurance and equity interest in trusts, annuities and managed investment accounts. Non financial assets include the primary residence, vehicles, investment real estate and business equity. Debt includes mortgage and home equity loans for primary residence and investment real estate, credit card balances and other loans.

The stock shares are defined as the fraction of financial wealth invested in stock. Financial wealth includes: Liquid accounts, bonds, stocks, mutual

funds, retirement accounts, the cash value of life insurance and a miscellaneous group of other assets. In turn liquid accounts include: Checking and savings accounts, money market deposits and mutual funds, saving funds not invested in stocks, certificates of deposits, call /cash accounts. The category of bonds include all types of bonds: Savings, government, tax exempt, mortgage backed, corporate and foreign bonds. The category of stocks includes directly held stocks. Stocks may be held also through mutual funds and the SCF reports these separately. Investment funds include: Stock mutual funds, combination mutual funds, other mutual funds. Retirement accounts are divided into IRA and Keoghs, job based 401k accounts, thrift savings accounts. The miscellaneous financial assets group is made of other types among else managed assets like personal annuities and trusts.

In the paper stocks include all stocks held directly and indirectly through mutual funds, IRAs, Keoghs and thrift type retirement accounts, or in other managed accounts like trusts and annuities. The Survey of Consumer Finance provides only a qualitative answer with regard to the fraction of a mutual fund, retirement account or managed account that is invested in stocks. We thus had to make an imputation to reach a quantitative figure. For mutual funds we imputed a fraction of 1 if the fund is defined as a stock mutual fund and of $\frac{1}{2}$ if it is a combination fund. For IRAs and Keoghs we attribute a fraction to stocks of 1 if the fund is mostly invested in stocks,

of $\frac{1}{2}$ if it is split between stocks and either bonds or money market assets and of $\frac{1}{3}$ if it is invested in all the three categories of assets. For stocks in the remaining categories of funds and managed accounts we attribute the full value to stock if it is described as mostly invested in stocks and $\frac{1}{2}$ if it is described as split.

Table 1: Conditional shares by wealth percentiles (Baseline)

Percentiles	0-25	25-50	50-75	75-100	Top 5%
Data (College)	56.2	60.9	62.7	67.6	63.3
Transaction cost (L)	66.3	91.8	95.9	81.8	69.4
Transaction cost (H)	60.0	60.3	73.1	74.2	69.0

Share of wealth invested in stocks for agents with positive holdings by wealth quartiles and separately for the top 5 percentiles of the wealth distribution. The second line reports the 2007 SCF data for the sub-sample of college graduates. The share is computed as fraction of stocks in financial wealth. The third and last line report the model-generated values in the case of low (L) and high (H) transaction cost. Simulated data use the college graduates' earnings process and the baseline parametrization with risk aversion parameter $\alpha = -4$.

[Footnote to table 2]

Share of wealth invested in stocks for agents with positive holdings by wealth quartiles and separately for the top 5 percentiles of the wealth distribution conditional on age. The first column defines the age groups. The top panel reports the 2007 SCF data for the sub-sample of college graduates. The share is computed as fraction of stocks in financial wealth. The second and third panel report the model-generated values in the case of low (L) and high (H) transaction cost. Simulated data use the college graduates' earnings process and the baseline parametrization with risk-aversion parameter $\alpha = -4$.

Table 2: Conditional shares by wealth percentiles and age (Baseline)

	Quart. I	Quart. II	Quart. III	Quart. IV	Top 5%
Data (College)					
20-30	43.6	53.6	45.2	56.7	72.1
30-40	54.3	63.5	61.9	61.6	63.4
40-50	59.7	68.8	65.2	68.6	59.1
50-60	67.2	69.4	69.4	68.8	62.8
60-70	51.4	59.3	70.4	67.7	58.9
70-80	39.6	41.5	61.7	68.4	68.7
TC low					
20-30	50.6	61.1	63.1	77.0	83.3
30-40	75.0	83.6	90.3	96.0	97.1
40-50	92.5	96.1	97.3	96.4	93.2
50-60	97.3	96.7	93.1	82.8	76.2
60-70	98.0	90.8	77.0	64.6	62.5
70-80	83.6	99.7	88.9	73.1	65.1
TC high					
20-30	7.7	47.1	46.7	55.8	57.7
30-40	57.5	59.1	54.7	58.9	73.9
40-50	53.9	62.1	75.7	75.7	70.4
50-60	70.0	76.5	79.1	69.9	69.8
60-70	73.5	76.7	75.1	68.5	70.6
70-80	56.2	70.1	75.6	72.3	66.7

[Locate footnote to table 2 here]

Table 3: Conditional shares by wealth percentiles (Sensitivity analysis on risk aversion)

Percentiles	0-25	25-50	50-75	75-100	Top 5%
Data (College)	56.2	60.9	62.7	67.6	63.3
Transaction cost (L)	56.9	86.2	94.1	98.4	98.8
Transaction cost (H)	69.1	55.1	71.2	89.7	92.5

Share of wealth invested in stocks for agents with positive holdings by wealth quartiles and separately for the top 5 percentiles of the wealth distribution. The second line reports the 2007 SCF data for the sub-sample of college graduates. The share is computed as fraction of stocks in financial wealth. The third and last line report the model-generated values in the case of low (L) and high (H) transaction cost. Simulated data use the college graduates' earnings process. Risk aversion is reduced by increasing α to -1 .

Table 4: Conditional shares by wealth percentiles (Sensitivity analysis on the return process, high cost scenario)

Percentiles	0-25	25-50	50-75	75-100	Top 5%
Data (College)	56.2	60.9	62.7	67.6	63.3
Positive correlation	59.3	63.2	70.7	70.4	65.4
Dividend	58.9	60.9	71.4	71.8	63.6

Share of wealth invested in stocks for agents with positive holdings by wealth quartiles and separately for the top 5 percentiles of the wealth distribution. The second line reports the 2007 SCF data for the sub-sample of college graduates. The share is computed as fraction of stocks in financial wealth. The third line reports the model-generated values in the case of positive correlation between stock returns and earnings shocks. The fourth line reports the model-generated values when the stock return is made by a separate dividend and price gain component. Simulated data use the college graduates' earnings process, the baseline parametrization and the high transaction cost.

Table 5: Conditional shares by wealth percentiles (Two education groups,high cost scenario)

Percentiles	0-25	25-50	50-75	75-100	Top 5%
Data	55.9	59.7	59.3	61.7	62.5
Model	57.8	58.2	71.2	75.5	72.8

Share of wealth invested in stocks for agents with positive holdings by wealth quartiles and separately for the top 5 percentiles of the wealth distribution. The second line reports the 2007 SCF data for the whole population. The share is computed as fraction of stocks in financial wealth. The third line reports the model-generated values using the baseline parametrization and the high transaction cost. The simulated population uses both agents with the college educated earnings process and agents with the high-school educated earnings process.

Table 6: Wealth-to-nonfinancial income by age and wealth: Two educational groups

Percentile	25	50	75	95
Data (Financial wealth)	0.07	0.45	1.51	5.38
Model	0.22	0.89	2.73	4.38

Age	20-30	30-40	40-50	50-60	60-70	70-80
Data (Financial wealth)	0.39	0.57	1.15	2.14	3.30	3.50
Model	0.54	0.79	1.76	3.09	4.60	2.87

The top panel reports the wealth-to-nonfinancial income ratio in the data and in the simulated population at the percentiles indicated in the top line of the panel. The bottom panel reports the wealth-to-nonfinancial income ratio in the data and in the model by ten-year age groups. The data refer to the whole population and the model is simulated using both college and high school educated agents.

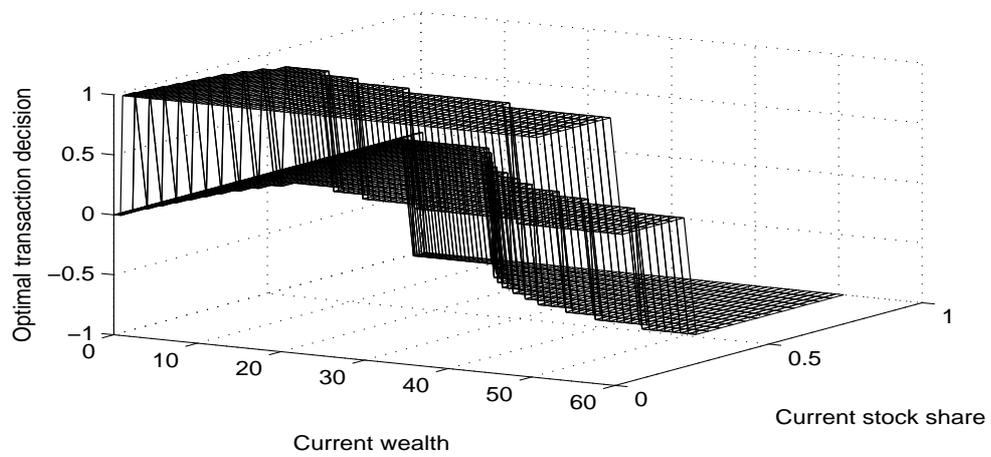


Figure 1: Transaction decision rule for a 45 year old, high school graduate endowed with the lowest earnings shock. On the vertical axis + 1 means “buy stock”, 0 means “no action”, -1 means “sell stock”. On the horizontal axis current wealth and current stock share are the two state variables.

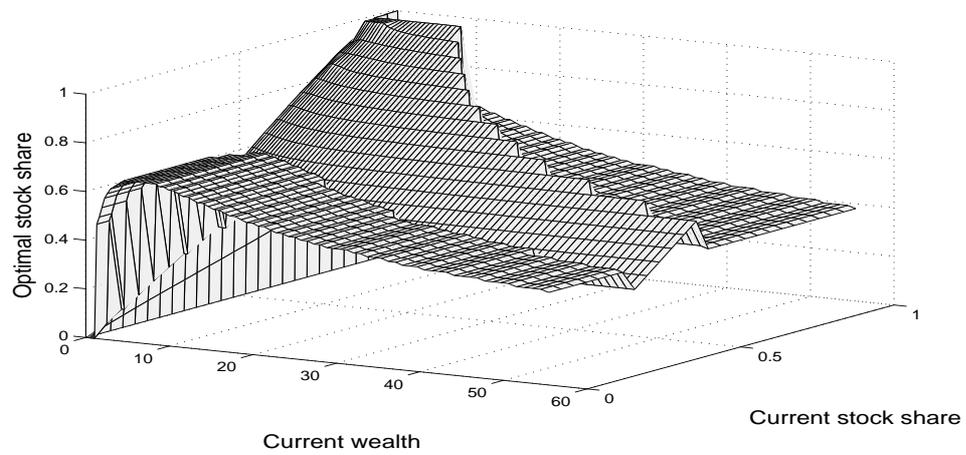


Figure 2: Stock share decision rule for a 45 year old, high school graduate endowed with the lowest earnings shock. On the vertical axis is the optimal stock share decision. On the horizontal axis current wealth and current stock share are the two state variables.

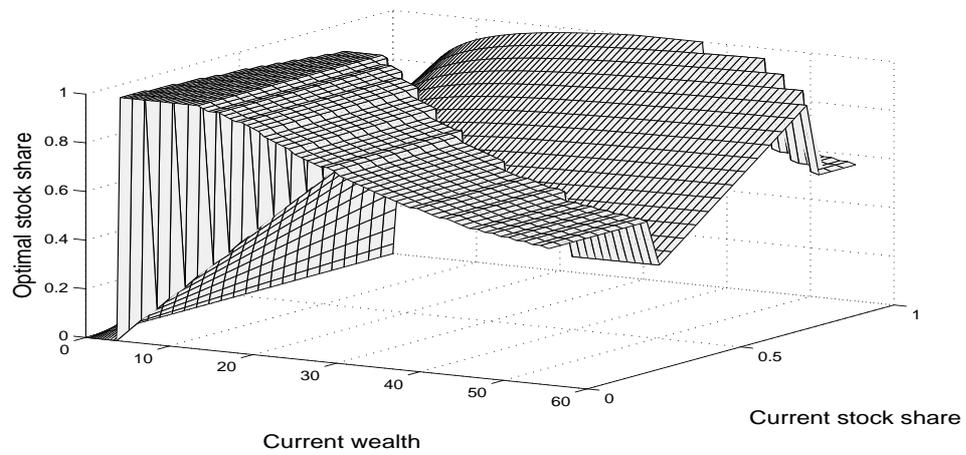


Figure 3: Stock share decision rule for a 45 year old college graduate endowed with the highest earnings shock. On the vertical axis is the optimal stock share decision. On the horizontal axis current wealth and current stock share are the two state variables.

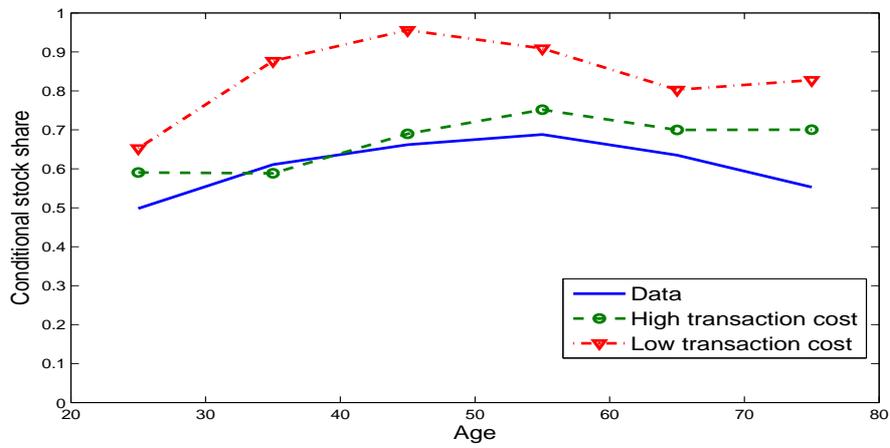


Figure 4: Life-cycle stock share for participants. Both data and model profiles refer to college graduates. The profiles in this figure are based on the baseline parametrization, with the coefficient that controls risk aversion α set to -4 .

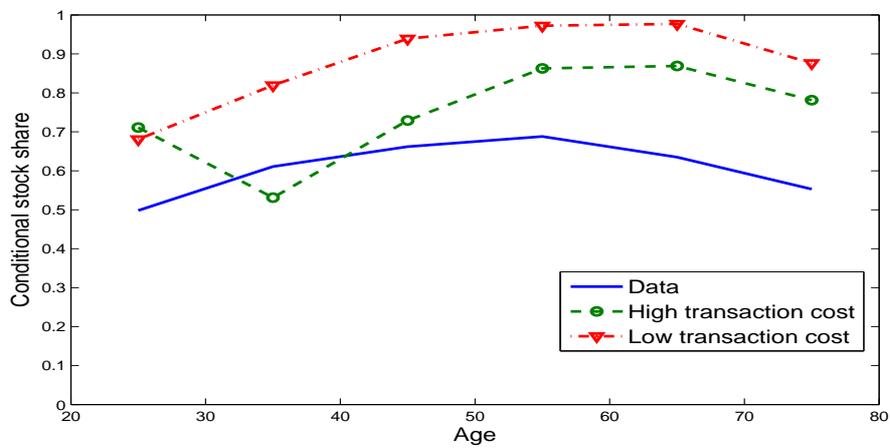


Figure 5: Life-cycle stock share for participants. The profiles in this figure report sensitivity analysis on risk aversion, with α set to -1 . Both data and model profiles refer to college graduates.

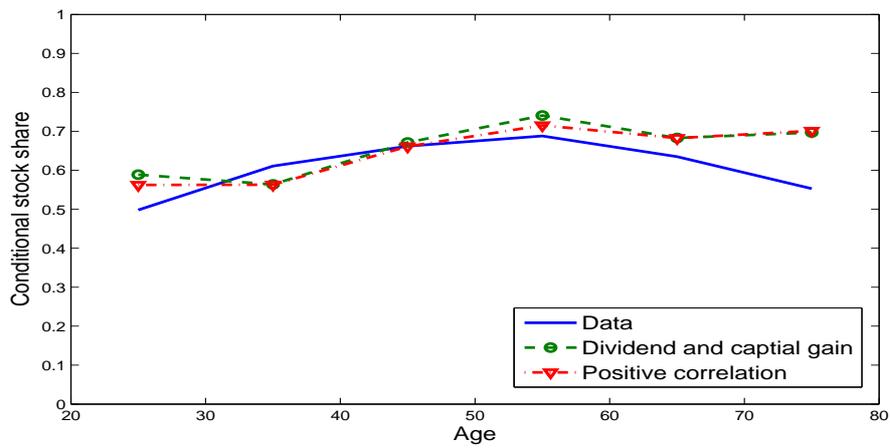


Figure 6: Life-cycle stock share for participants: Sensitivity analysis on the return process, high cost scenario. Both data and model profiles refer to college graduates. Two cases are considered: Positive correlation between stock returns and earnings, and a return process with separate dividend and capital gain.

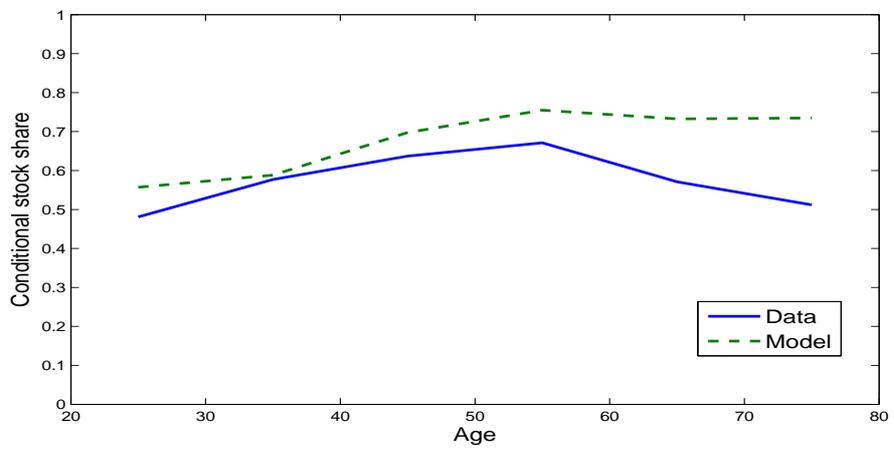


Figure 7: Life-cycle stock share for participants: two educational groups, baseline parametrization, high cost scenario.