Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures

Annette Vissing-Jørgensen[1]

Department of Economics, University of Chicago

vissing@uchicago.edu

First draft: May 4, 1998

This draft: March 2, 1999

[1]This paper is a substantially revised version of Chapter 2 of my MIT thesis, 1998. I thank Olivier Blanchard, Ricardo Caballero, Jim Heckman, Lars Hansen, and the participants in the workshop in empirical economics at the University of Chicago for helpful comments.

Abstract

The paper uses micro data on income and asset holdings from the Panel Study of Income Dynamics to analyze reasons for nonparticipation and for heterogeneity in portfolio choice within the set of stock market participants. The focus of the paper is on non-financial income and cost of participating in the stock market.

I find evidence of a strong positive effect of mean non-financial income on the probability of stock market participation and on the proportion of wealth invested in stocks conditional on being a participant. The volatility of non-financial income is found to have a strong negative impact on these two quantities. Both these results are consistent with the theoretical literature on portfolio choice in the presence of non-financial income. However, only a small or insignificant effect of the covariance of non-financial income with the stock market return on portfolio choice is present.

Four different costs of stock market participation are considered, an entry cost, a fixed transactions cost, a proportional transactions cost, and a per period participation cost. The first three of these costs lead to structural state dependence in the stock market participation decision and in the proportion of financial wealth invested in stocks. A dynamic sample selection model shows evidence of strong state dependence and thus economically important entry and transactions costs. The per period participation cost does not lead to structural state dependence, but a censored regression model with unobservable stochastic threshold is estimated which allows for heterogeneity in the per period partipation cost and provides estimates of its distribution in the cross-section. I estimate the median per period participation cost to be around $100-$200.

1 Introduction

Households differ dramatically in their portfolio choices. Among households surveyed in the wealth supplements to the Panel Study of Income Dynamics, 49.66 percent of those with positive financial wealth do not hold positions in either stocks nor bonds[1]. An additional 23.04 percent hold stocks but no bonds, whereas 14.79 percent hold bonds but no stocks. Only 12.50 percent hold both stocks and bonds. Furthermore, even within the set of households who hold both stocks and bonds, observed portfolio compositions differ substantially. These facts are not consistent with predictions from standard models of agents maximizing expected lifetime utility subject to initial wealth and the possibility of investing in all existing assets at zero transaction costs.

In the standard model of lifetime consumption and portfolio choice of Samuelson (1969) and Merton (1969), (1971), agents live off income generated by their invested wealth. The set of available assets includes a conditionally riskless asset and N risky assets. Without transaction costs this model predicts that agents should take positions in all existing assets counter to the frequently observed zero holdings. The optimal portfolio of risky assets and the split between risky and riskless assets will vary across agents with different preferences, wealth and investment horizon. Conditions on return distributions/utility functions have been derived, under which differences in wealth and investment horizon across agents should not lead to differences in portfolio choice. Investment horizons are irrelevant if agents face a constant investment opportunity set (i.i.d returns). CRRA preferences are sufficient for wealth not to matter. Heterogeneity in risk aversion always implies differences in portfolio choice.

It is well known that actual returns are not i.i.d implying potential heterogeneity in portfolio choices across age groups. There is less agreement as to whether CRRA utility is a reasonable approximation. Conditional on this assumption, evidence is starting to accumulate documenting heterogeneity in the (constant) coefficient of relative risk aversion. In a very interesting paper, Barsky et al (1997) document heterogeneity in risk aversion based on micro data from the Health and Retirement Study. About 12000 respondents answered questions concerning gambles over lifetime income. The answers reveal considerable heterogeneity in risk tolerance and the survey measure of risk tolerance significantly predicts portfolio shares in stocks, bonds, treasury bills

[1] Averages of numbers from 1984, 1989 and 1994. Stockholdings include stocks held through mutual funds and in IRAs. The bond measure does not include money market bonds, or Treasury bills. Some non-bond categories of wealth are included in the bonds measure (life insurance policies, collections, and rights in trusts or estates).

and checking and savings accounts with the expected signs. However, the incremental predictive power of risk tolerance is low and the remaining unexplained variation in portfolio choice large, even after including wealth, income and demographic controls (including age).

The purpose of this paper is to determine empirically whether accounting for differences in nonfinancial income patterns across households and for costs of participating in the stock market helps explain the remaining heterogeneity in observed portfolio choices.

The theoretical literature on non-financial income and background risk predicts three effects. A larger mean of non-financial income should lead agents to invest a larger fraction of financial wealth in stocks, since agents with alternative sources of income can rely on this for consumption purposes should their financial investments fail. The variance of non-financial income should have a negative effect on the proportion invested in stocks due to background risk inducing more risk averse behavior. A non-zero covariance of non-financial income with stock returns should cause a hedging component of asset demand of the opposite sign of the covariance.

I focus on the two largest sources on nonfinancial income, namely labor income and income from privately held businesses. I concentrate mainly on the decision to hold stocks and the proportion of financial wealth held in stocks. Future work will also consider bonds. I use income data from the PSID, 1979-1993, and three observations of portfolio choice from the 1984, 1989 and 1994 wealth supplements. A two step procedure is used, similar to the one followed in previous papers, in which the first step consist of estimating the relevant moments of income processes which are then used as regressors in the second step focusing on portfolio choice. My results based on probit and tobit regressions document economically important and statistically significant mean and variance effects of non-financial income on portfolio choice. Weaker evidence is found for a covariance effect. The results concerning the mean and variance effect confirms the findings of Guiso et al. (1996) using a somewhat different methodology and a different data set.

Four different costs of stock market participation are considered, an entry cost, a fixed transactions cost, a proportional transactions cost, and a per period participation cost. The first three of these costs lead to structural state dependence in the stock market participation decision and in the proportion of financial wealth invested in stocks. A dynamic sample selection model shows evidence of strong state dependence and thus economically important entry and transactions costs. The per period participation cost does not lead to structural state dependence, but a censored regression model with unobservable stochastic threshold is estimated which allows for heterogeneity in the per period participation cost and provides estimates of its distribution in the cross-section. I estimate the median per period participation cost to be around \$100-\$200.

Aside from its importance for understanding portfolio choice and thus the determination of prices of financial assets, the question of whether observed heterogeneity is consistent with optimizing behavior has important policy implications. One set of implications concerns the optimal portfolio composition of a social security trust fund. If the observed heterogeneity in portfolio choice represents a rational response to different economic conditions, then imposing the same portfolio composition on all households could imply large welfare losses. Another policy implication concerns the effect of taxation of labor/business income on agents' lifetime utility. A proportional tax will have the effect of decreasing the mean and standard deviation of after-tax labor/business income which will affect agents' consumption and portfolio choice. This effect of taxation has been analyzed by Elmendorph and Kimball (1991). The results documenting strong mean- and variance-effects of non-financial income on portfolio choice, emphasize the importance of this issue.

2 Related literature

The recent asset pricing literature has paid much attention to the effects of labor income on portfolio choice and general equilibrium asset pricing, especially focusing on the effects of uninsurable idiosyncratic shocks. It has been known since Merton(1971) that the existence of certain non-financial income should cause agents with HARA utility to invest a larger fraction of their financial wealth in risky assets. This is the 'mean-effect' mentioned above. However, labor income is generally risky. With an incomplete set of financial markets agents cannot rely solely on financial markets to insure themselves. Furthermore, moral hazard problems prevent insurance contracts between labor income earners and potential insurers. The uninsurable part of nonfinancial income implies a 'variance-effect' on portfolio choice. Gollier and Pratt (1996) consider the effect of unfair background risks, i.e. risks with nonpositive expectations. In a one period model they show that all familiar DARA utility functions are risk vulnerable, meaning that any unfair background risk makes risk-averse agents behave in a more risk averse way. Viceira (1997) extends this result to a multi-period model in which wealth accumulation is endogenous. Viceira

(1997) furthermore clearly shows the effect on portfolio choice of the sign of the covariance of labor income innovations with the stock return. Positive covariance of the stock return and innovations to permanent or transitory income generates a negative hedging component of asset demands and vice versa for negative covariance[2]. The theory of background risk can be applied to uninsurable income from privately held businesses as well as to labor income.

Empirical work testing these predictions are still at a quite early stage. Let me briefly mention three papers, two which focuses on labor income risk and one which also considers business risk. To estimate household level income processes a (long) panel data set of income observations and at least one observation of portfolio choice is needed. Guiso, Japelli et al. circumvent the need for a panel data set by using the 1989 Bank of Italy Survey of Household Income and Wealth. In this survey respondents are asked to distribute probability weights to given intervals of inflations and nominal labor and pension income changes one year ahead. Based on the answers an estimate of expected income variance can be constructed (but the covariance-effect cannot be tested). This variable is an economically and statistically significant predictor of the proportion of wealth in risky assets. The level of income enters positively, which could be interpreted as support for the mean-effect predicted by theory. Gakidis (1997) uses 7 years of income data from the PSID and the 1984 wealth supplement. He estimates labor income processes by demographic groups defined by occupation, age, and education as an alternative to estimating income processes by household. The most important finding is a significantly negative effect of the probability of zero income events on the probability of being a stockholder and on the proportion held in stocks conditional on being a stockholder. He finds no evidence of a variance-effect aside from the zero-income probabilities, but some evidence of a mean effect. The covariance effect is not considered, probably due to the short sample. Heaton and Lucas (1997) use the 1979-1990 Panel of Individual Tax Return Data and exploits the panel dimension to calculate, by household, the standard deviation of labor income and business income and the covariance of these two income components with the S&P500 stock return. These are then used as regressors in a random effects regression with the proportion of financial wealth invested in stocks as the dependent variable. The results are inconclusive most likely due to the poor

[2]Several papers analyze the general equilibrium effects of uninsurable income risk on asset prices. See, for example, Mankiw (1986), Heaton and Lucas (1996), Constantinides and Duffie (1996), Krusell and Smith (1997), and Telmer et al. (1998).

quality of asset data in this data set[3]. When the sum of labor and business income is used they find evidence of a positive variance-effect (counter to the prediction from theory). However, when labor and business income are included separately, and the sample restricted to those with average business income above \$500 the standard deviation of business income has the expected negative sign and is significant. None of the covariance variables are significant.

Concerning the part of the paper which attempts to identify stock market participation cost structures from a panel of portfolio choice data, the most related paper is Mulligan and Sala-i-Martin (1996). They consider the demand for money versus interest bearing assets and use the 1983-89 panel from the Survey of Consumer Finances. The 1983 value of financial asset enters significantly with a positive sign in a probit model for holding interest bearing assets in 1989. The authors interpret this as evidence of a positive start-up cost of investing in interest bearing assets. Mulligan and Sala-i-Martin furthermore estimate the cost of participating in the market for interest bearing assets to be between \$50 and \$200.

Other related papers on adjustment costs and asset pricing are Luttmer (1997) and Marshall and Parekh (1998), both focusing on costs of adjusting consumption.

3 Framework

3.1 Costs of stock market participation

Consider the optimization problem of a household which maximizes expected lifetime utility given an exogenous stream of nonfinancial income and faced with the opportunity to invest in two assets, a risky asset and a conditionally riskless asset. The risky asset represents the stock market. The riskless asset is a catchall for less risky financial assets such as T-bills, bank accounts etc. Below I will refer to it as T-bills.

This optimization problem is standard, except that I consider the following three costs of investing in the stock market.

F^E : Stock market entry cost. This represents the time/money spent understanding what the stock market is and determining the household's optimal mix between stocks and T-bills. Before the invention of low cost mutual finds, F^E would include the cost of learning how to buy a well diversified portfolio (the actual trading costs will be included in the cost F^T

[3] Asset holdings must be estimated based on information on dividends, interest income and capital gains.

discussed below). Currently, since one can buy the stock market index through a mutual fund, the main part of F^E is likely to be the cost of time spent determining the household's optimal portfolio shares for stocks versus bonds. Add to that the cost of time spent setting up accounts.

F^P : Per period stock market participation cost. F^P represents time spent dealing with stock market investments. In their study of the demand for money, Mulligan and Sala-i-Martin (1996) mention as a cost of holding interest bearing assets the extra time spent doing taxes (schedule D). To the extent that households feel it necessary to follow the stock market if they invest in it, the time spent doing this would also be included in F^P. With time varying conditional asset return distributions, the theory of dynamic hedging suggests that households should be following the stock market in order to form more precise expectations of future returns and change their portfolios accordingly. From a less theoretical perspective, the increased number of TV programs about the stock market during recent years - a time with large increases in stock market participation - suggests that households do indeed spend time following the market.

F^T : Trading cost. This cost is likely to have two components. Firstly, a fixed part F_0^T, representing fixed commissions and the value of time spent trading. Secondly, a variable part $F_1^T \left|N_{it} - N_{i,t-1}\right|$, representing proportional transactions costs, e.g. the bid-ask spread and the variable part of commissions. N_{it} denotes the number of stocks held by household i between date t and date $t+1$.

3.2 Value function and implied structure of policy functions

The value function for the optimization problem of a household that faces the investment opportunities outlined above is as follows.

$$V_{it}\left(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}\right) = \max_{\{0,1\}} \left\{V_0\left(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}\right), V_1\left(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}\right)\right\}$$

where

$$V_0\left(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}\right) = \max_{C_{it}} \left\{U\left(C_{it}\right) + \delta E V_{i,t+1}\left(0, 0, W_{it+1}, S_{it+1}\right)\right\}$$

$$s.t. \quad W_{i,t+1} = \left(1 + r_{f,t}\right)\left(W_{it} - I\left(Y_{i,t-1} = 1\right)\left(F_0^T + F_1^T N_{i,t-1}\right)\right) + \omega_{it} - C_{it}$$

$$V_1\left(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}\right) = \max_{C_{it}, N_{it}} \left\{U\left(C_{it}\right) + \delta E V_{i,t+1}\left(N_{it}, 1, W_{it+1}, S_{it+1}\right)\right\}$$

$$\begin{aligned} s.t. \quad W_{i,t+1} &= \left(1 + r_{f,t} + \alpha_{it}\left(r_{s,t} - r_{f,t}\right)\right)\left(W_{it} - I\left(Y_{i,t-1} = 0\right) F^E \right. \\ &\quad \left. - I\left(N_{i,t-1} \neq N_{it}, Y_{i,t-1} = 1\right)\left(F_0^T + F_1^T \left|N_{it} - N_{i,t-1}\right|\right) - F^P\right) + \omega_{it} - C_{it} \end{aligned}$$

$$\alpha_{it} = \frac{P_{st} N_{it}}{W_{it} - I\left(Y_{i,t-1} = 0\right) F^E - I\left(N_{i,t-1} \neq N_{it}, Y_{i,t-1} = 1\right)\left(F_0^T + F_1^T \left|N_{it} - N_{i,t-1}\right|\right) - F^P}$$

$V^1(.)$ is the expected lifetime utility if the household chooses to participate in the stock market in the current period. $V^0(.)$ is the corresponding quantity if the household chooses not to participate in the current period. $U(.)$ is the per period utility function and $\delta < 1$ the time discount factor. The state variables at date t are:

$N_{i,t-1}$: The number of stocks owned in the previous period, if any

$Y_{i,t-1}$: The lagged participation decision. $Y_{i,t-1}$ equals 1 if the household participated in the stock market in the previous period and 0 otherwise

W_{it} : Financial wealth of household i at t

S_{it} : State variables characterizing the process for nonfinancial income for household i at t

As for the asset returns, $r_{f,t+1}$ denotes the net return on T-bills, and $r_{s,t+1}$ the net return on stocks before paying entry/participation/transactions costs. α_{it} refers to the proportion of W_{it}, net of costs, which the household invests in stocks.

$V_0(.)$ is standard aside from the term $I\left(Y_{i,t-1} = 1\right)\left(F_0^T + F_1^T N_{i,t-1}\right)$ which is the cost of exiting the stock market. The indicator function $I(.)$ equals 1 if the argument is true and 0 otherwise. In the expression for $V_1(.)$, the entry cost F^E must be paid if the household was not in the market in the previous period. The trading cost $F_0^T + F_1^T \left|N_{it} - N_{i,t-1}\right|$ must be paid if the household participated in the previous period but chooses to change its stockholdings. The per period participation cost F^P similarly represents a reduction in wealth available for investment.

Whether the household is assumed to be finitely or infinitely lived, a closed form solution to the above problem is in general not obtainable. Even in the absence of costs of stock market participation, the theory on portfolio choice in the presence of non-financial income does not allow a closed form solution to be derived when non-financial income is risky (and less than perfectly correlated with the stock return) and/or investment opportunities vary over time. Progress has been made by Campbell and Viceira (1996) for the case of time varying investment opportunities and the previously mentioned paper by Viceira (1997) for the model with uninsurable labor income. By log-linearizing the Euler-equations and the budget constraints they obtain analytical solutions to approximate problems. Unfortunately, the log-linearization constants are complicated functions of the underlying parameters and numerical solutions must still be used to determine, for example, the effect of an increase in the standard deviation of permanent income growth on the optimal portfolio share of the risky asset.

However, although a closed form solution is not in general available, the structure of the policy functions is clear from setup. The policy function for the participation decision, Y_{it}, will depend on:

$N_{i,t-1}$ (with a positive sign) if $F_1^T > 0$. The data set I use below does not contain information about the number of stocks held, N_{it} only about the dollar value of stock held, $P_t^s N_{it} = \alpha_{it} W_{it}^{disp}$ and thus the fraction of wealth invested in stocks, $\alpha_{it} = N_{it} \frac{P_{st}}{W_{it}^{disp}}$ where $W_{it}^{disp} \equiv W_{it} - I\left(Y_{i,t-1} = 0\right) F^E - I\left(N_{i,t-1} \neq N_{it}, Y_{i,t-1} = 1\right)\left(F_0^T + F_1^T \left|N_{it} - N_{i,t-1}\right|\right) - F^P$. I therefore include $\alpha_{i,t-1}$ rather that $N_{i,t-1}$.

$Y_{i,t-1}$ (with a positive sign) if $F^E > 0$, or $F_0^T > 0$. $F^E > 0$ makes it more likely that a household will participate in the current period if it participated in the previous period since the entry cost is already paid. In the above I have assumed that the entry cost depreciates fully upon exit. If, alternatively, F^E did not have to be repaid upon reentry, the term $I\left(Y_{i,t-1} = 0\right) F^E$ in the expression for $V_1(.)$ would be replaced by $I\left(\max\left\{Y_{i,t-1}, Y_{i,t-2}, ...\right\} = 0\right) F^E$. The variable $\max\left\{Y_{i,t-1}, Y_{i,t-2}, ...\right\}$ would then be the relevant explanatory variable in the binary choice model rather than $Y_{i,t-1}$. Intermediate cases with partial depreciation will lead to positive but gradually smaller coefficients on still higher lags of the participation decision. By acting as an exit cost, a positive value for F_0^T similarly implies that participation at $t-1$ makes participation at t more likely.

W_{it} (with a positive sign) if $F^E > 0$ or $F^P > 0$ or $F_0^T > 0$, and in addition the optimal dollar amount invested is increasing in wealth. This would be the case for utility functions of the DARA type.

S_{it}

Age$_{it}$ of head of household: With finite lifetimes the investment horizon depends on age. If investment opportunities are time varying, optimal portfolio choice depends on the investment horizon and thus on age. The shape of the age-dependence will depend on the specific processes for non-financial income and asset returns, so age squared is also included in the relation to allow a more flexible functional form. Age could also matter for other reasons, for example because of health risks being age dependent.

Education$_{it}$ of head of household: Education could be significant over and above determining labor income if F^E, F^P or F^T (or returns obtained) depends on education or if preferences differ across education groups

The policy function for the number of stocks held, N_{it}, conditional on participation implies a policy function for α_{it}, the proportion of financial wealth invested in stocks. This policy function will have the following arguments:

$N_{i,t-1}$ (with a positive sign) if $F_0^T > 0$ or $F_1^T > 0$, since $N_{it} = N_{it-1}$ for households who participated both last period and the current period but do not find it optimal to trade. Again, since I do not observe $N_{i,t-1}$ I instead include $\alpha_{i,t-1}$ in the model for α_{it}.

$Y_{i,t-1}$ if $F^E > 0$ or $F^T > 0$ and wealth is a significant determinant of α. However, since F^E and F^T are likely to be small fractions of wealth for most of the households who chose to participate, the effect of $Y_{i,t-1}$ on α_{it} (conditional on α_{it} being positive) is likely to be negligible. The empirical results confirms this, and I therefore focus on estimations which exclude $Y_{i,t-1}$ in the equation for α_{it}.

W_{it}, S_{it}, Age$_{it}$, Education$_{it}$

Thus, by estimating a reduced form model of Y_{it} and α_{it} and determining if the lagged policy variables are significant with effects of the expected signs, it is possible to determine if the costs F^E, F_0^T, and F_1^T are large enough to significantly affect portfolio choice. In terms of distinguishing between these three types of costs, it is possible to draw conclusions about F_1^T from the effect of $\alpha_{i,t-1}$ on $P\left(Y_{it} = 1\right)$. A significant effect of $\alpha_{i,t-1}$ on α_{it} is due to trading costs, but which type (fixed or proportional) cannot be separately identified. A significant effect of $Y_{i,t-1}$ on $P\left(Y_{it} = 1\right)$ is due to either entry costs or fixed trading costs ($F^E > 0$ or $F_0^T > 0$).

The table below summarizes the implications of F^E, F_0^T, F_1^T for the lagged policy variables in the policy functions.

Policy function	State variable	Significant if	Expected sign
Y_{it}	$Y_{i,t-1}$	$F^E > 0$ or $F_0^T > 0$	$+$
Y_{it}	$N_{i,t-1}$ (or $\alpha_{i,t-1}$)	$F_1^T > 0$	$+$
α_{it}	$N_{i,t-1}$ (or $\alpha_{i,t-1}$)	$F_0^T > 0$ or $F_1^T > 0$	$+$

As for the per period participation cost F^P it does not lead to structural state dependence.[4] The last section of the paper considers in more detail what can be learned about F^P from reduced form modelling.

4 Data

For the purpose of estimating household level income processes a fairly long panel of income information is need. This motivates the use of the Panel Study of Income Dynamics for my analysis, along with the availability of several years of wealth and portfolio information in this data set. I use the Survey Research Center sample of the PSID which was representative of the civilian noninstitutional population of the US when the study was started in 1968. The PSID has tracked all original family units and their adult offspring over time, so with low attrition rates the sample remains representative as long as offsprings are included. I excluded the poverty sample and the Latino sample.

The last year for which final release data are available is 1993. From the 1968-93 family files I construct a data set containing information for each of the households ever in the sample during this period. I use the family files rather than the individual files since wealth information is available at the household level. There are 6322 such households (after excluding the poverty and latino samples). For split-offs, information for years prior to the split-off was coded as missing.

Wealth information from the 1984, 1989 and 1994 supplements is used to calculate net financial wealth, defined as the sum of cash (checking and savings accounts, money market

[4]'Structural state dependence' refers to a situation in which past decisions affect current decisions by changing the nature of the current optimization problem. More on this later.

bonds, Treasury bills, including such assets held in IRA's), bonds (bond funds, cash value in life insurance policies, collections, rights in trusts or estates), and stocks (shares of stock in publicly held corporations, mutual funds, or investment trusts, including stocks in IRA's). To identify entries for which imputations were used, I use the wealth information as given in the family files instead of the wealth supplement files. Imputed values for cash, bonds or stocks can then be coded as missing. Topcoding of wealth or income variables is very rare in the PSID and topcoded variables were left at their topcodes.

Although nothing prevents households having a portfolio share for a given asset above one, the PSID wealth data does not allow one to observe this due to the way the wealth questions are formulated. For example, the questions asked concerning stock holdings are "Do you (or anyone in your family living there) have any shares of stock in publicly held corporations, mutual funds, or investment trusts, including stocks in IRA's?" and "If you sold all that and paid off anything you owed on it, how much would you have?". Thus, a household who had borrowed to invest more than its total financial wealth in stocks would be recorded as having a portfolio share for stocks of one. Similarly, it is not possible to identify negative portfolio shares (short sales) because negative values of stocks, bonds and cash are coded as zeros in the PSID.

I define nonfinancial income as labor income plus business income of both head and spouse. Income and wealth variables are deflated by the consumer price index, with 1982-84 as basis year. Household years in which the head is a student or in which the head is older than 80 years are dropped. It was also necessary to drop the households with the three largest values of wealth in order to be able to calculate some of the estimators to follow.

Estimates of each of the three income moments $\mu_{it} \equiv E_t\left(Y_{i,t+1}\right), \sigma^2_{it} \equiv \left(V_t(Y_{i,t+1})\right)^{1/2}$, and $cov_{it} \equiv cov_t(Y_{it+1}, R^s_{t+1})$ are then constructed at the household level. One based on the 5-year window around 1984[5], one based on the 5-year window around 1989, and one based on the 15-year period 1978-1992. For the 5-year windows only households with no changes in head or spouse during any of the two windows are included. For the 15-year window only households with no changes in head or spouse during this period are used. For the 5-year windows, households with 3 or more non-zero observations or more of labor plus business/farm income are used. For the 15-year window households with 10 or more observations are used. If nonfinancial income

[5] This correspond to interview years 1983-87, since income for the previous year is reported when a household is interviewed.

zero in a particular year, that value is not used to calculate the moments. If nonfinancial income is zero in 1984 or 1989, the three moments are set to zero for that household in that year. The stock return used for calculating the covariance of stocks and nonfinancial income is the real value weighted NYSE index.

The use of windows is motivated by potential time-variation in nonfinancial income. This time-variation can, if present, be used to construct IV estimators which takes measurement error in the income moment estimates into account. The time-variation also allows fixed-effect estimators to be used although the small number of portfolio observations causes some problem in this regard. I will return to this issue later.

Since the latest available income information refers to 1992, it is not possible to construct a window around 1994. The 1994 portfolio information is therefore only used for descriptive statistics and for the regressions with the 15 year window. I do not calculate the covariance of income components with the real stock market return for the 5-year windows. It is unlikely that the covariance of an income component and the stock market return can be estimated to any level of precision with 5 years of data, and aside from that, it is not clear that this covariance would change much over time should we be able to estimate it precisely. The reason for using a 15 year period in stead of the entire sample for each household for the covariance estimation is that many households change composition over time, even if only changes in heads and spouses are considered. Therefore, restricting the sample to households with the same head and spouse for all years would imply a very small (and far from representative) sample.

The base line sample with data based on the 5-year windows around 1984 and 1989 contains 1092 households.

5 Results

5.1 Basic facts about heterogeneity in portfolio choice

For comparison with previous studies, Table 1 confirms for the present data set, the well known fact that in any given year only a fraction of households with positive financial wealth participates in the stock market or in the bond market. An upward trend in stock market participation is clear from the PSID data. Of households with positive financial wealth 44.06 percent participated in the stock market in 1994, up from 34.12 percent in 1989 and 28.47 percent in 1984. Within the set of stockholders, both the median and mean of stockholdings in dollars and then mean percentage of financial wealth held in stocks increases strongly between 1989 and 1994[6]. To give a representative picture which can provide information about the US population as a whole, these numbers are based on all households with positive financial wealth in the PSID[7].

With three observations of portfolio choice for a group of households, it is possible to analyze patterns of participation and trading over time. Many papers have emphasized widespread non-participation in one or both of these markets based on a cross-section of households. By following households over time it is possible to determine whether households either stay in/out of a given market or whether there is widespread movements in and out of markets.

The results are shown in Fig. 1-3 and Table 2, all based on households with positive financial wealth[8]. Fig. 1a. focuses on the set of households for which portfolio information is available for both 1989 and 1994. The figure plots the 1994 share of financial wealth held in stocks against the 1989 share. In the absence of costs of participating in the stock market, and with no nonfinancial income and i.i.d asset returns, standard finance theory predicts that all households should be at a point along the 45 degree line in this figure (the origin not included). With nonfinancial income and/or returns which are not i.i.d points off the 45 degree line but in the interior of the first quadrant are potentially consistent with theory. Only if there are costs of participating in the stock market can the large number of households at the origin or along one of the axes (71.72 percent) be explained. Previous evidence based on cross sections of households would lead us to expect many observations of zero stockholdings in each year. Somewhat surprisingly, the figure shows that many households participate in the stock market in one year but not the other. These are the points along the axes forming an angle in the graph. 28.10 percent of households are on this angle, not including the origin. Fig. 1b. shows similar results based on the 1984-89 panel. Notice that many of the points on the angles are far from zero. This reflects households who move from a zero to a substantially positive fraction of wealth in stocks or the other way around and for which the entry/exit thus does not correspond to 'marginal' changes in stockholdings as a percentage of financial wealth.

Fig. 1c. focuses on households with positive financial wealth for which three observations

[6] This is not a necessary consequence of the stock market boom, since with more participants each participant could in theory hold the same amount or the same percent of financial wealth in stocks in 1994 as in 1989.

[7] Thus not all of the households used for the tabulations are the same for all three years due to split-offs etc.

[8] As for Table 1 I do not drop households with changes in household composition.

of portfolio choice are available. The change in the share of financial wealth held in stocks between 1989 and 1994 is plotted against the 1984-89 change. A 'triplet' of lines is apparent. The vertical line corresponds to households who did not participate in the stock market in 1984 or 1989 but did participate in 1994 (points showing zero change between 1984 and 1989 are all for non-participant who had a zero share in both years). The horizontal line corresponds to households who participated in 1984 but not in 1989 or 1994. The most interesting line is the downward sloping one, which shows that many households entered the stock market some time between 1984 and 1989, but left the market again some time between 1989 and 1994[9]. It is tempting to interpret this as households who entered but got scared by the market crash in October 1987. However, Fig. 2c. shows a similar pattern for bond holdings. A more plausible explanation for this pattern is large changes in optimal portfolio shares combined with a fixed per period cost of participating in the market. The graphs for the remaining component of financial wealth, cash, are shown in Fig. 3. The lines are the 'reverse' of those shown in Fig.1 and 2, which is intuitive.

5.2 Evidence on the importance of nonfinancial income for portfolio choice

5.2.1 Main results

The results from estimating a sample selection model of the stock market participation decision and of the proportion of financial wealth invested in stocks, are overall very encouraging in terms of documenting a mean-effect and a variance-effect of nonfinancial income. Some, but weaker, evidence is found of a covariance-effect.

Table 3 shows the results for the sample selection model, estimated on the base line sample of 1092 households using portfolio data from 1984 and 1989. There is clear evidence of a positive mean-effect and a negative variance-effect of nonfinancial income on the probability of stock market participation. The effects are economically important as well as statistically significant. The marginal effects (evaluated at the means of the right hand side variables) show that an increase in mean real nonfinancial income of 10000 dollars (in 1982-84 prices) increases the probability of participation by about 4 percentage points. The effect of a change in the standard deviation of real nonfinancial income is of a similar magnitude but negative. In the equation

[9] From the available information it is not possible to determine if households entered and then left the market again, or left and then reentered, in years in between 1984 and 1989 or in between 1989 and 1994.

for the optimal proportion of financial wealth invested in stocks, the effect of a higher standard deviation of nonfinancial income is negative and significant. An increase by \$10000 reduces the optimal share of wealth invested in stocks by about 4 percentage points. The mean of nonfinancial income also enters with the expected sign, but is not significant.

Education enters the regressions with the same signs as has been found in other studies. Households the head of which is more educated are more likely to hold stocks and to hold a large proportion of their wealth in stocks, if any. If interaction terms of the education dummies and the lagged participation decision/lagged α are included (not shown), they are generally insignificant. This indicated that the significance of education is more likely to reflect a correlation between education and preferences than between education and the costs of stock market participation.

The net effect of an increase in real financial wealth on the probability of stock market participation and on the proportion invested in stocks is positive at all wealth levels observed in the sample. This could be entirely due to the presence of costs of investing in the stock market, and does not necessarily mean that the coefficient of relative risk aversion is decreasing in wealth.

Since $\widehat{\mu}_{it}$ and $\widehat{\sigma}_{it}^2$ are likely to be noisy estimates of the true parameters μ_{it} and σ_{it}^2, a measurement error problem may be present in the results shown in Table 3. To determine if this is the case, I reestimate the model using instrumental variables techniques. I instrument for $\widehat{\mu}_{it}$ and $\widehat{\sigma}_{it}^2$ using as instruments the lagged values $\widehat{\mu}_{it-1}$ and $\widehat{\sigma}_{it-1}^2$. The IV estimator for the probit model is the one described in Amemiya (1978). For the proportion of financial wealth invested in stocks I use the estimator from Lee (1981). The results are shown in Table 4. In the probit, the effect of mean nonfinancial income is about twice as large as in Table 3 while the effect of the standard deviation of nonfinancial income is about six times larger. This clearly shows the importance of instrumenting to avoid downward bias in the estimates due to measurement error. For the proportion of wealth invested in stocks, the effects of nonfinancial income is again much larger in the IV estimations and mean nonfinancial income is now significant.

Turning to the covariance effect, Table 5 shows the results based on defining the income variables on the 15-year window 1979-1993 and using portfolio data for 1984, 1989, and 1994. The covariance of nonfinancial income with the stock return enters both the probit and the equation for the proportion in stock with the expected sign, but is not significant. If the correlation is used in stead of the covariance (not shown), it enters significantly in the equation for the proportion in stocks with a coefficient of -0.17. The 10th percentile for the correlation (in the cross section of households) is -0.20. The 90th percentile equals 0.23. Thus even a move from the 10th to the 90th percentile in terms of correlation of nonfinancial income with the stock return would only change the proportion invested in stocks by 7 percentage points. Based on this it is not surprising that the covariance (and the correlation if that is used instead) is insignificant in the probit model.

In sum, the results strongly support the theoretical prediction of a positive effect of mean nonfinancial income on the probability of stockholding and the optimal share of financial wealth invested in stocks, and of a negative effect of the variance of nonfinancial income. Some evidence is found of a negative covariance effect but this effect is both economically and statistically weaker.

5.2.2 Accounting for individual effects

In dynamic panel data models, even individual effects which are not correlated with regressors other than the lags can lead to inconsistent estimates of all parameters. Consider, for example, a dynamic probit model as the one considered presently. If the initial conditions of the process could be assumed truly exogenous, or the process could be assumed to be in equilibrium, consistent estimates of the parameters could be obtained using maximum likelihood estimation. In most cases, neither of these assumptions are plausible. In that case one could consider a fixed effects approach. This does, however, only leads to consistent estimates as $T \to \infty$. In short samples the bias of the fixed effects estimator can be very large as shown in Heckman (1981). Heckman instead proposes an approximate random effects estimator which is still biased but much less so than the fixed effects estimator. This approximate random effects estimator assumes that the individual effects are uncorrelated with the regressors other than the lags. For the panel data probit model with one lag, Heckman's estimator can be used for $T = 2$ or larger. Recently, Honore and Kyriazidou (1998) have shown that provided that 4 or more observations per individual are available, one can consistently estimate the parameters of panel data logit models with individual effects which are allowed to be correlated with the explanatory variables. Since I only have $T = 2$ (or $T = 3$ if only one 15 year window is used to calculate the income variables), I cannot use this estimator. I will return to Heckman's estimator in the next section.

If the model did not include lags one could account for individual effects which are correlated with the regressors in the binary choice model of stock market participation by estimating a conditional logit model rather than a probit. Individual (household) effects are likely to be correlated with the income regressors if the fixed effect represents heterogeneity in preference parameters. Consider a case in which we are trying to determine the effect of the mean and standard deviation of labor and business income on the probability of being a stockholder. Suppose that households are heterogeneous in terms of their coefficient of relative risk aversion but that we do not have household level risk aversion measures. Less risk averse households are more likely to be stockholders. However, they are also more likely to self-select into riskier jobs or to become business owners and will therefore tend to have higher standard deviation of labor (business) income, and also higher mean labor (business) income to the extent the risk is compensated by a higher mean. Thus the regressors will be endogenous and the coefficient estimates on both the mean and the standard deviation will be upward biased. The effect of mean income will be exaggerated and we may get an unexpected positive or insignificant coefficient for labor (business) income risk on stockholdings (similar problems arise for the covariance-effect as for the variance-effect). The Survey of Consumer Finances contains a self-reported measure of risk aversion which several papers have found significant in regressions involving stockholdings[10], confirming the findings of Barsky et al. (1997). The latter also documented an economically large although not statistically significant of risk tolerance on the probability of being self-employed.

To partially address this issue I chose to estimate a conditional logit model even though lags cannot be included when $T < 4$, and the results therefore should be interpreted with some caution. The idea of this fixed-effects model is to condition on the sum of the discrete dependent variable over the sample (at the household level). In the resulting conditional likelihood the individual effects cancel, see Chamberlain (1980).

The estimation results are shown in Table 7. Households who do not enter or exit the stock market in one of the two years contribute zeros to the conditional likelihood and are dropped.

[10]See Blume and Zeldes (1994) and Bertaut and Haliassos (1995). Due to the lack of time dimension the SCF by itself cannot be used to test for the three effects of non-financial income on portfolio choice. One could consider estimating the moments of income by demographic groups using the PSID and then use these as regressors in an analysis based on SCF asset data and the SCF risk aversion measure.

The 1996 PSID for the first time includes a measure of risk aversion similar to the one from the Health and Retirement Study used by Barsky et al. I will return to the potential use of this measure.

This leaves 526 households. As an alternative sample I therefore drop the restriction on no household composition changes. This results in a larger sample of 866 households, but with lower quality of the income data in the sense that a time series of income observation for a given household may not refer to the same head and spouse for all years. I furthermore consider a cutoff of 5 percent of financial wealth in stocks for being considered a stockholder. For three of the four resulting cases the mean-effect of nonfinancial wealth is significant, now for the probability of entering or exiting the stock market. There is no evidence of a variance effect. Financial wealth is significant when the larger sample is used. Given the assumed lack of time-variation in the covariance of nonfinancial wealth with the stock return, it is not possible to test whether the (weak) evidence of a covariance effect is robust to controlling for individual effects. Overall the conditional logit results show that the mean-effect is robust to controlling for individual effects by identifying the effect off time variation in stock market participation for each household, rather than from cross-sectional differences as done in the estimations above. The lack of evidence of a variance effect based on the present estimates does not necessarily mean that one is not present once controlling for individual effects. It is possible that there is insufficient time-variation in the standard deviation of nonfinancial income or that the time-variations are too small to be precisely estimated.

An alternative approach to controlling for individual effects is possible if one is willing to make the assumption that risk aversion (and no other preference parameters) is the main component of the individual effect. For the first time, the 1996 questionnaire of the PSID includes a series of questions designed to provide an estimate of relative risk aversion (or equivalently the inverse of relative risk aversion called risk tolerance). The methodology used is the same as the one described in the paper by Barsky et al. (1997). The risk aversion questions are only asked to employed respondents. By using the household identification numbers in the 1994-96 early release files the estimates of risk tolerance can be merged back into the sample. I only merge backwards to 1989 to avoid serious selection problems due to fact that a household has to remain in the sample until 1996 and the head has to be employed for risk tolerance estimates to be available. The risk tolerance estimate can then be included in an estimation for the 1989 cross-section as an estimate of the part of the individual effect which is correlated with the regressors. As the conditional logit model, this approach does not allow for lags of the participation decision.

The results given in Table 6 show that the risk tolerance measure is significant in the probit for stock market participation at the 10 percent level and has the expected positive sign. However, the clear evidence of a mean and variance effect of nonfinancial income remains. This is comforting given the mixed results of conditional logit estimations.

5.3 Evidence on the structure of stock market participation costs

5.3.1 Main results

The coefficient on the lagged participation decision is positive and highly significant in the probit in both Table 3 and 4. As outlined in the theoretical section this leads to the conclusion that $F^E > 0$ or $F_0^T > 0$. Although I am not able to quantify the dollar amounts of these costs, their economic importance can be considered. The discrete change in the probability of being a stockholder in period t caused by being a stockholder in period $t-1$ is 30 percentage points based on the results from the IV estimations in Table 4.

The coefficient on $\alpha_{i,t-1}$ is also positive and significant in the probit model. A 10 percentage point increase in $\alpha_{i,t-1}$ increases the probability of being a stockholder in period t by about 2 percentage points, again based on the results from Table 4. Thus $F_1^T > 0$ and F_1^T is large enough to be economically important.

Turning to the proportion invested in stocks, α_{it}, the positive coefficient on $\alpha_{i,t-1}$ signifies that $F_0^T > 0$ or $F_1^T > 0$. As for the size of the effect, Table 4 shows that a 10 percentage point increase in $\alpha_{i,t-1}$ increases $\alpha_{i,t}$ by about 5 percentage points, conditional on participation.

5.3.2 Accounting for individual effects

As discussed above, the significance of lags could be due to unaccounted for serially correlated individual effects rather than structural state dependence. To determine if this is the case I estimated the dynamic probit model using the estimator of Heckman (1981). The model estimated and the sample used is the same as in Table 3, aside from α_{it-1} not being included. The next version of the paper with include α_{it-1} as well as Y_{it-1}. The coefficient on Y_{it-1} remains essentially unchanged (and the mean and standard deviation of nonfinancial income remains significant). This indicates that the significance of Y_{it-1} in the probit is indeed evidence that $F^E > 0$ or $F_0^T > 0$, and not due to unobserved heterogeneity. Heckman's estimator for the probit model can be generalized to the sample selection model. Derivations and results for this will also be included in the next version of the paper.

6 Inference about the size of $\mathbf{F}^P$

Let $C_{it}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it})$ and $\alpha_{it}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}) = N_{it}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}) \frac{P_{st}}{W_{it}^{disp}}$ be the solutions to the optimization problem for a stockholder (the same problem considered earlier)

$$\begin{aligned} V_1(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}) &= \max_{C_{it}, N_{it}} \{U(C_{it}) + \delta E V_{i,t+1}(N_{it}, 1, W_{it+1}, S_{it+1})\} \\ s.t. \quad W_{i,t+1} &= (1 + r_{f,t} + \alpha_{it}(r_{s,t} - r_{f,t})) W_{it}^{disp} + \omega_{it} - C_{it} \end{aligned}$$

with W_{it}^{disp} defined as before. Define the certainty equivalent excess return on stocks over T-bills as the nonstochastic excess return which would make the household indifferent between the stochastic portfolio return $1 + r_{f,t} + \alpha_{it}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it})(r_{s,t} - r_{f,t})$ and the certain portfolio return $1 + r_{f,t} + \alpha_{it}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}) r_{it}^{ce}$, holding the policy functions $C_{it}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it})$ and $\alpha_{it}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it})$ fixed. The certainty equivalent excess return for household i, period t, will in general depend on the state variables $N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}$ since households with a high value of $\alpha_{it}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it})$ are likely to have a high certainty equivalent return. This implies a positive correlation between α_{it} and r_{it}^{ce} in the cross-section.

Using the above definition of r_{it}^{ce}, the optimality condition for stock market participation for household i in period t can be rewritten from

$$\text{Participate if} \quad V_1(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}) \geq V_0(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it})$$

to

$$\text{Participate if} \quad \alpha_{it}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}) W_{it}^{disp} r_{it}^{ce}(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}) \geq F_{it}^P.$$

The above participation condition is similar to the condition given by Mulligan and Sala-i-Martin (1996), except that I do not assume the certainty equivalent return to be constant across households. Notice that F_{it}^P is allowed to differ across households. The objective of the analysis to follow is to estimate the cross sectional distribution of F^P in the population. To be able to do this one needs additional assumptions about

a) The value of α_{it} $\forall i$, or a model of it, $\alpha_{it}\left(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}\right)$.

b) The value of r_{it}^{ce} $\forall i$, or a model of it, $r_{it}^{ce}\left(N_{i,t-1}, Y_{i,t-1}, W_{it}, S_{it}\right)$.

c) The correlation of F_{it}^{P} with the other variables in the model

I first consider the simple benchmark case of homogeneous α_{it} and r_{it}^{ce} and then move to the more realistic case of heterogeneity.

6.1 Case 1: Homogeneous α_{it} and r_{it}^{ce}

Assume that

a) $\alpha_{it} = 1$ $\forall i$

b) $r_{it}^{ce} = 0.04$ $\forall i$

c) Wealth and participation costs are uncorrelated in the cross section.

Given assumptions a) and b) the stock market participation condition reduces to:

$$\text{Participate if} \quad W_{it}^{disp} 0.04 \geq F_{it}^{P}.$$

Since the incentive to participate is linear in wealth one can estimate the cross sectional distribution of F_{it}^{P} directly from the wealth distribution. For example, if 27 percent of households with wealth of $10000 participates, then 27 percent of these households must have had participation cost below $0.04 * 10000 = \$400$. Given assumption c) this implies that 27 percent of all households must have had participation costs below $400. By splitting the sample into 10 wealth deciles and using this approach for each decile, one obtains 10 estimates of points on the cumulative distribution function for the cross sectional distribution of F_{it}^{P}.

Results:

The result is shown in Figure 4 for the sample of all households with positive financial wealth and assuming $\alpha = 1$. The median per period participation cost is around $600 (real 1982-84 dollars) for 1994, higher for 1984 and 1989. Since even among very rich households not all hold stock, the estimated CDF does not reach 1 at any wealth level[11]. Figure 5 shows the same calculation under the assumption that α_{it} for each household equals the fraction of stock market wealth in total financial wealth of the sample for year t. This leads to estimates of the median per period participation cost of about $350 for each year. Interpreting the per period participation cost as the cost of additional time spent doing taxes and spent following the market, this seems

[11] The point corresponding to the last wealth decile is not included in the graph.

somewhat high. However, as shown in the next section more careful modelling brings down the estimate of the median to about \$100-\$200.

6.2 Case 2: Heterogeneous α_{it} and r_{it}^{ce}. A censored regression model with unobservable stochastic threshold

Assume that α_{it} is determined by the model

$$\alpha_{it} = \begin{cases} \alpha_{it}^{*} = \exp\left(x_{it}'\beta + u_{it}\right) \text{ if } \alpha_{it}^{*} W_{it}^{disp} r_{it}^{ce} \geq F_{it}^{P} \\ 0 \end{cases}$$

$$\ln F_{it}^{P} = \mu_{\ln F^P} + \eta_{it}$$

$$\ln r_{it}^{ce} = x_{it}'\varphi + v_{it}$$

$$\Rightarrow \ln F_{it}^{P} - \ln r_{it}^{ce} = \underbrace{\left(\mu_{\ln F^P} - x_{it}'\varphi\right)}_{\equiv x_{it}'\overline{\varphi}} + \underbrace{\left(\eta_{it} - v_{it}\right)}_{\equiv \varepsilon_{it}} = x_{it}'\overline{\varphi} + \varepsilon_{it} \tag{1}$$

$$\begin{bmatrix} u_{it} \\ \varepsilon_{it} \end{bmatrix} | x_{it}, W_{it}^{disp} \sim N_2 \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_u^2 & \sigma_{u\varepsilon} \\ \sigma_{u\varepsilon} & \sigma_\varepsilon^2 \end{bmatrix} \right)$$

where x_{it} is a vector containing the arguments of the policy function.

A few comments are in order. Firstly, this model will lead to a conservative (large) estimate of the distribution of the per period participation cost, since the optimal proportion invested in stocks (conditional on participation) is assumed positive for all households. Had I (as earlier) used a simple linear functional form, $\alpha_{it}^{*} = x_{it}'\beta + u_{it}$, the resulting participation cost distribution would have been centered at much lower values of costs. If a household with large wealth does not participate in the stock market, the linear specification could attribute this to a negative value of α_{it}^{*} and a short sales constraint, rather than a large value of F_{it}^{P}. Secondly, the model allows for F^E, F_0^T and F_1^T (as well as F_{it}^{P}) since these are subtracted in W_{it}^{disp}.[12] Thirdly, F_{it}^{P} is assumed to not depend on x_{it}. This is clearly an approximation, since I suggested that a large part of F_{it}^{P} is likely to be a cost of time, and since x_{it} contains variables which describes labor income.

[12] Assuming that the observed portfolio at t is post-trade, the wealth data are for W_{it}^{disp}, not for W_{it}.

For $\alpha_{it} = 0$ this model implies

$$\begin{aligned}
P\left(\alpha_{it}|x_{it}, W_{it}^{disp}\right) &= P\left(\ln\left(\alpha_{it}^{*}W_{it}^{disp}\right) \leq \ln F_{it}^{P} - \ln r_{it}^{ce}|x_{it}, W_{it}^{disp}\right) \\
&= P\left(x_{it}'\beta + u_{it} + \ln W_{it}^{disp} \leq x_{it}'\overline{\varphi} + \varepsilon_{it}|x_{it}, W_{it}^{disp}\right) \\
&= P\left(u_{it} - \varepsilon_{it} \leq x_{it}'\overline{\varphi} - x_{it}'\beta - \ln W_{it}^{disp}|x_{it}, W_{it}^{disp}\right) \\
&= \Phi\left(\frac{x_{it}'\overline{\varphi} - x_{it}'\beta - \ln W_{it}^{disp}}{\sqrt{\sigma_u^2 + \sigma_\varepsilon^2 - 2\sigma_{u\varepsilon}}}\right)
\end{aligned}$$

whereas for $\alpha_{it} = \alpha_{it}^{*}$

$$\begin{aligned}
P\left(\alpha_{it}|x_{it}, W_{it}^{disp}\right) &= P\left(\ln\left(\alpha_{it}^{*}W_{it}^{disp}\right) \geq \ln F_{it}^{P} - \ln r_{it}^{ce}|x_{it}, W_{it}^{disp}\right) \\
&= P\left(x_{it}'\beta + u_{it} + \ln W_{it}^{disp} \geq x_{it}'\overline{\varphi} + \varepsilon_{it}|x_{it}, W_{it}^{disp}\right) \\
&= P\left(\varepsilon_{it} \leq -x_{it}'\overline{\varphi} + x_{it}'\beta + u_{it} + \ln W_{it}^{disp}|u_{it}, x_{it}, W_{it}^{disp}\right) P\left(u_{it}|x_{it}, W_{it}^{disp}\right) \\
&= \Phi\left(\frac{-x_{it}'\overline{\varphi} + \ln\left(\alpha_{it}^{*}W_{it}^{disp}\right) - \frac{\sigma_{u\varepsilon}}{\sigma_u^2}u_{it}}{\sqrt{\sigma_\varepsilon^2\left(1 - \left(\frac{\sigma_{u\varepsilon}}{\sigma_u\sigma_\varepsilon}\right)^2\right)}}\right)\frac{1}{\sqrt{2\pi\sigma_u^2}}\exp\left(-\frac{1}{2}\left(\frac{u_{it}}{\sigma_u}\right)^2\right)
\end{aligned}$$

where $u_{it} = \ln(\alpha_{it}^{*}) - x_{it}'\beta$ is observed since α_{it}^{*} is.

Using these formulas, the parameters β, $\overline{\varphi}$, σ_u^2, σ_ε^2, $\sigma_{u\varepsilon}$ can be estimated by maximum likelihood estimation using the data for all households i in a given period t. What can we infer about the distribution of F^P based on this? Below I outline ways of estimating $\mu_{\ln F^P}$ and $\sigma_{\ln F^P}^2$. By symmetry of the normal distribution, the estimate of $\mu_{\ln F^P}$ is also an estimate of the median of $\ln F^P$. An estimate of the median of F^P is then given by

$$\widehat{\text{Median}}\left(F^P\right) = \exp\left(\widehat{\mu_{\ln F^P}}\right).$$

Furthermore, if $\ln F_{it}^P = \mu_{\ln F^P} + \eta_{it}$ is normally distributed, F_{it}^P will be lognormal with

$$\begin{aligned}
\mu_{F^P} &\equiv E\left(F_{it}^P\right) = \exp\left(\mu_{\ln F^P} + \frac{1}{2}\sigma_{\ln F^P}^2\right) \\
\sigma_{F^P}^2 &\equiv V\left(F_{it}^P\right) = \exp\left(2\mu_{\ln F^P} + \sigma_{\ln F^P}^2\right)\left[\exp\left(\sigma_{\ln F^P}^2\right) - 1\right].
\end{aligned}$$

Inference about $\mu_{\ln F^P}$:

If the first element in x_{it} is a constant term, then

$$\begin{aligned} x_{it}'\overline{\varphi} &= \mu_{\ln F^P} - x_{it}'\varphi \\ \text{i.e. } \overline{\varphi}_1 &= \mu_{\ln F^P} - \varphi_1 \quad \Leftrightarrow \mu_{\ln F^P} = \overline{\varphi}_1 + \varphi_1 \\ \overline{\varphi}_{-1} &= -\varphi_{-1}. \end{aligned}$$

Subscript 1 refers to element 1 of the vector in question, and subscript -1 to the remaining elements. To estimate $\mu_{\ln F^P} = \overline{\varphi}_1 + \varphi_1$ we need not only the MLE estimate of $\overline{\varphi}_1$ but also an estimate of φ_1.

It was assumed above that $\ln r_{it}^{ce} = x_{it}'\varphi + v_{it}$. Under the additional assumption

$$E\left(v_{it}|x_{it}, W_{it}^{disp}\right) = 0$$

we have

$$E\left(\ln r_{it}^{ce}|x_{it}, W_{it}^{disp}\right) = x_{it}'\varphi = \varphi_1 + (x_{it})_{-1}'\varphi_{-1}$$

and thus

$$\begin{aligned} E\left(\ln r_{it}^{ce}\right) &= E\left(x_{it}\right)'\varphi = \varphi_1 + \left(E\left(x_{it}\right)_{-1}\right)'\varphi_{-1} = \varphi_1 + \left(E\left(x_{it}\right)_{-1}\right)'\overline{\varphi}_{-1} \\ &\Leftrightarrow \ \varphi_1 = E\left(\ln r_{it}^{ce}\right) - \left(E\left(x_{it}\right)_{-1}\right)'\overline{\varphi}_{-1} \\ &\Rightarrow \ \mu_{\ln F^P} = \overline{\varphi}_1 + \varphi_1 = \overline{\varphi}_1 + E\left(\ln r_{it}^{ce}\right) - \left(E\left(x_{it}\right)_{-1}\right)'\overline{\varphi}_{-1} \end{aligned}$$

Using the MLE estimate of $\overline{\varphi}_{-1}$ and the cross-sectional means of each variable in $(x_{it})_{-1}$ this equation can be used to estimate φ_1, provided we are willing to make an assumption about the unconditional cross sectional mean $E\left(\ln r_{it}^{ce}\right)$. I assume that

$$r^{ce} \sim \text{lognormal}\left(\mu_{r^{ce}}, \sigma_{r^{ce}}^2\right)$$

which implies

$$\ln r^{ce} \sim N\left(\ln\left(\mu_{r^{ce}}^2\right) - \frac{1}{2}\ln\left(\mu_{r^{ce}}^2 + \sigma_{r^{ce}}^2\right), \ln\left(1 + \frac{\mu_{r^{ce}}^2}{\sigma_{r^{ce}}^2}\right)\right).$$

I use $\mu_{r^{ce}} = 0.03$, $\sigma_{r^{ce}}^2 = 0.01^2$ as benchmark values, and then consider a couple of alternative parameter choices to determine how sensitive the resulting estimate of $\mu_{\ln F^P}$ is to the choice.

Inference about $\sigma^2_{\ln F^P}$:

The maximum likelihood estimation provides an estimate of

$$V\left(\varepsilon_{it}|x_{it}, W_{it}^{disp}\right) = \sigma_\varepsilon^2 = \sigma_\eta^2 + \sigma_v^2 - 2\sigma_{\eta v}.$$

It seems plausible to assume a zero covariance between the error term for the participation cost and the error term for the certainty equivalent excess return on stocks, $\sigma_{\eta v} = 0$. Then

$$\sigma^2_{\ln F^P} = \sigma_\eta^2 = \sigma_\varepsilon^2 - \sigma_v^2.$$

The maximum likelihood estimate of σ_ε^2 is therefore an estimate of the largest possible value of $\sigma^2_{\ln F^P}$ (corresponding to $\sigma_v^2 = 0$). As a simple alternative, I also calculate the results under the assumption that $\sigma_v^2 = \frac{1}{2}\sigma_\varepsilon^2$.

Results:

Table 8 shows the median, mean and standard deviation of the cross-sectional distribution of the per period participation cost. The results are based on maximum likelihood estimation of the censored regression model with unobservable stochastic threshold as outline above. The model is estimated using the 1984 and 1989 data for the base line sample of 1092 households. The variables included in x_{it} are the same as those in Table 3. The median of F_{it}^P is \$103 (real 1982-84 dollars) when it is assumed that $\mu_{r^{ce}} = 0.03$, $\sigma^2_{r^{ce}} = 0.01^2$, $\sigma_v^2 = 0$. This number is not affected by the chosen value of $\sigma_v^2 = 0$. A median per period participation cost of \$103 dollars seems plausible, as it corresponds to about 10 hours of time when time is valued at \$10. Increasing $\mu_{r^{ce}}$ to 0.05, increases the estimate of the median of F_{it}^P to \$178. Lowering $\sigma^2_{r^{ce}} = 0.01^2$ only affects the estimated of the medean per period participation cost slightly.

However, assuming that all of the conditional variance of $\varepsilon_{it} = \eta_{it} - v_{it}$ is due to η_{it} (the error term for the log-participation cost) leads to implausibly large values for the mean and standard deviation of F_{it}^P. Under the alternative assumption that half of the conditional variance of ε_{it} is due to v_{it}, the estimates of the mean and standard deviation of F_{it}^P are much lower, but still seem implausible. As discussed in the context of the case 1 estimate of the CDF, this is due to the fact that even among very rich households not everyone holds stocks.

In sum, the results concerning the per period participation cost suggests that the PSID portfolio choice data are consistent with a median per period participation cost of about \$100-\$200. Due to the fact that some very rich households do not hold stocks, the variance of the cross-sectional distribution is implausibly large. Adding costs of stock market participation to the standard portfolio choice model cannot the behavior of these rich households in the sense that the required costs would have to be implausibly large.

7 Conclusion

Observed portfolio choices are not consistent with standard finance theory in the absence of a fixed cost of entering or staying in the stock market. Many households do not participate in the stock market at any point in time even when attention is limited to households with positive financial wealth. Furthermore, a substantial number of households move in or out of the stock market (and/or the bond market) over time. The present paper has emphasized nonfinancial income and costs of stock market participation as part of the explanation of the observed heterogeneity in portfolio choices.

I have argued elsewhere (Vissing-Jorgensen (1998)) that the nonparticipation phenomenon should be considered an important part of the solution to the equity premium puzzle. This is the case if the consumption growth of nonstockholders covaries substantially less with the stock return than the consumption growth of stockholders. Empirical evidence based consumption data from the Consumer Expenditure Survey confirmed that this was the indeed the case. This indicates that the primary reason for nonparticipation is not that nonstockholders are faced with nonfinancial income which is highly correlated with the stock market return. The findings of the present paper confirm this. Only a small or insignificant effect of the covariance of nonfinancial income with stock returns on participation and portfolio choice is found. Rather there is strong evidence of a mean-effect and a variance-effect.

Determining the reasons for heterogeneous portfolio choice and stock market nonparticipation is crucial, not only for having confidence in those results concerning the equity premium puzzle. Stock market participation has increased dramatically during the postwar period. The positive results concerning the contribution of limited stock market participation to the solution of the equity premium puzzle suggests that this may have had substantial effects on asset prices. To analyze this issue it is essential to understand the main reasons for nonparticipation.

The analysis of the present paper of the nature and size of costs of stock market participation has parallels to the literature on investment under adjustment costs. Caballero, Engel and Haltiwanger (1995) emphasize how an average adjustment rate function which is increasing in mandated investments can give rise to time-varying sensitivity of aggregate investment to aggregate shocks. In a similar way, the sensitivity of the stock price to e.g. a shock to nonfinancial income in a general equilibrium model with costs of stock market participation, is likely to differ depending of the cross-sectional distribution of households along the state variables. For example, the number of household who chose to enter the stock market or to change the number of stocks held in response to a shock to nonfinancial income, will depend on how many households are close to the point where it becomes worthwhile to adjust. Developing these parallels may contribute towards understanding issues regarding time varying stock market liquidity and trading volume.

References

[1] Amemiya, T., 1978, "The Estimation of a Simultaneous Equation Generalized Probit Model", Econometrica, 46, 5, 1193-1205.

[2] Allen, F. and D. Gale, 1994, "Limited Market Participation and Volatility of Asset Prices", American Economic Review, 84, 4, 933-955.

[3] Arrondel, L. and A. Masson, 1990, "Hypothese du cycle de vie, diversification et composition du patrimoine: France 1986", Annales D'Economie et de Statistique, 17, 1-45.

[4] Avery, R.B. and G.E. Elliehausen, "Financial Characteristics of High-Income Families", Federal Reserve Bulletin, March 1996.

[5] Barsky, R. B., F. Thomas Juster, M. S. Kimball, and M. D. Shapiro, 1997, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study", Quarterly Journal of Economics, 2, 537-579.

[6] Basak, S. and D. Cuoco, 1997, "An Equilibrium Model with Restricted Stock Market Participation", Rodney L. White Center for Financial Research, Working Paper 001-97.

[7] Blume, M. E. and S. P. Zeldes, 1994, "Household Stockownership Patterns and Aggregate Asset Pricing Theories", working paper, Wharton.

[8] Caballero, R. J., E. M. R. A. Engel, and J. C. Haltiwanger, 1995, "Plant-Level Adjustment and Aggregate Investment Dynamics", Brookings Papers on Economic Activity, 2, 1-54.

[9] Campbell, J.Y., 1997, "Asset Prices, Consumption and the Business Cycle", paper prepared from the Handbook on Macroeconomics, edited by J.B. Taylor and M. Woodford.

[10] Campbell, J.Y. and J.H. Cochrane, 1995, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", NBER Working Paper No. 4995.

[11] Campbell, J. Y. and L. M. Viceira, 1996, "Consumption and Portfolio Decisions When Expected Returns are Time Varying", working paper, Department of Economics, Harvard University.

[12] Chamberlain, G., 1980, "Analysis of Covariance with Qualitative Data", Review of Economic Studies, 47, 225-238.

[13] Chamberlain, G., 1984, "Panel Data", Handbook of Econometrics, Vol. II, Chapter 22, Elsevier Science Publishers BV.

[14] Constantinides G.M. and D. Duffie, 1996, "Asset pricing with Heterogeneous Consumers, Journal of Political Economy, 104, 219-240.

[15] Dumas, B., 1989, "Two-Person Dynamic Equilibrium in the Capital Market", The Review of Financial Studies, 2, 2, 157-188.

[16] Elmendorph, D. W. and M. S. Kimball, 1991, "Taxation of Labor Income and the Demand for Risky Assets", NBER Working Paper No. 3904.

[17] Gollier, C. and J. W. Pratt, 1996, "Risk vulnerability and the tempering effect of background risk", Econometrica, 64, 5, 1109-1123.

[18] Grossman, S.J., and G. Laroque, 1990, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods", Econometrica, Vol. 58, No.1.

[19] Guiso, L., T. Jappelli, and D. Terlizzese, 1996, "Income Risk, Borrowing Constraints, and Portfolio Choice", American Economic Review, 86, 1, 158-172.

[20] Haliassos, M. and C.C. Bertaut, 1995, "Why do so few hold stocks?", Economic Journal, 1110-1129.

[21] Heaton, J. and D. Lucas, 1995, "The importance of investor heterogeneity and financial market imperfections for the behavior of asset prices", Carnegie-Rochester Ser. Public Pol., 1-38.

[22] Heaton, J. and D. Lucas, 1996, "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing", Journal of Political Economy, 104, 3.

[23] Heaton, J. and D. Lucas, 1997, "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk", Working Paper No. 234, Kellogg Graduate School of Management, Northwestern University.

[24] Heckman, J. J., 1981, "The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process" in "Structural Analysis of Discrete Data with Econometric Applications", edited by C. F. Manski and D. McFadden, Cambridge, MIT Press.

[25] Honore, B. E. and E. Kyriazidou, 1998, "Panel Data Discrete Choice Models With Lagged Dependent Variables", working paper.

[26] Jakubson, G., 1988, "The Sensitivity of Labor-Supply Parameter Estimates to Unobserved Individual Effects: Fixed- and Random-Effects Estimates in a Nonlinear Model Using Panel Data", 6, 3, 302-329.

[27] Krusell, P, and A. A. Smith, 1997, "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns", Macroeconomics Dynamics, 1, 387-422.

[28] Lee, Lung-Fei, 1981, "Simultaneous Equations Models with Discrete and Censored Dependent Variables", in "Structural Analysis of Discrete Data with Econometric Applications", edited by C. F. Manski and D. McFadden, Cambridge, MIT Press.

[29] Luttmer, E. G. J., 1997, "What level of Fixed Costs Can Reconcile Asset Returns and Consumption Choices", working paper, Northwestern University

[30] Mankiw, N. G., 1986, "The Equity Premium Puzzle and the Concentration of Aggregate Shocks", Journal of Financial Economics, 17, 211-219.

[31] Mankiw, N.G. and S. Zeldes, 1991, "The consumption of Stockholders and Non-Stockholders," Journal of Financial Economics, 17, 211-219.

[32] Marshall, D. A. and N. G. Parekh, 1998, "Can Costs of Consumption Adjustment Explain Asset Pricing Puzzles", forthcoming, Journal of Finance.

[33] Merton, R. C., 1969, "Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case", Review of Economics and Statistics, 51, 247-257.

[34] Merton, R. C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model", Journal of Economic Theory, 3, 373-413.

[35] Merton, R.C., 1987, "A Simple Model of Capital Market Equilibrium with Incomplete Information", Journal of Finance, 42, 3, 483-510.

[36] Merton, R. C., 1992, "Continuous-Time Finance", Blackwell Publishers, Cambridge, MA.

[37] Mulligan, C. B., and X. X. Sala-i-Martin, 1996, "Adoption of Financial Technologies: Implications for money demand and monetary policy", NBER Working Paper No. 5504.

[38] Poterba, J. M and A. A. Samwick, 1995, "Stock Ownership Patterns, Stock Market Fluctuations, and Consumption", Brookings Papers on Economic Activity, 1995:2, 295-372.

[39] Samuelson, P. A., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming", Review of Economics and Statistics, 51, 39-46.

[40] Sandmo, A., 1969, "Capital Risk, Consumption, and Portfolio Choice", Econometrica, 37, 4, 586-599.

[41] Sandmo, A., 1970, "The Effect of Uncertainty on Saving Decisions", Review of Economic Studies, 37, 353-360.

[42] Storesletten, K., C. Telmer, and A. Yaron, 1998, "Persistent Idiosyncratic Shocks and Incomplete Markets", working paper, Stockholm University.

[43] Viceira, L. M., 1997, "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income", working paper, Department of Economics, Harvard University.

[44] Vissing-Jorgensen, A., 1998, "Limited Stock Market Participation", Ph.D thesis, Department of Economics, MIT.

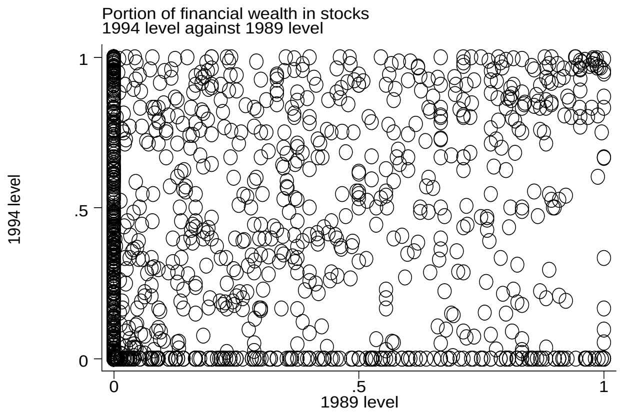

Fig.1a The fixed cost 'stock angle', 1989 and 1994 levels

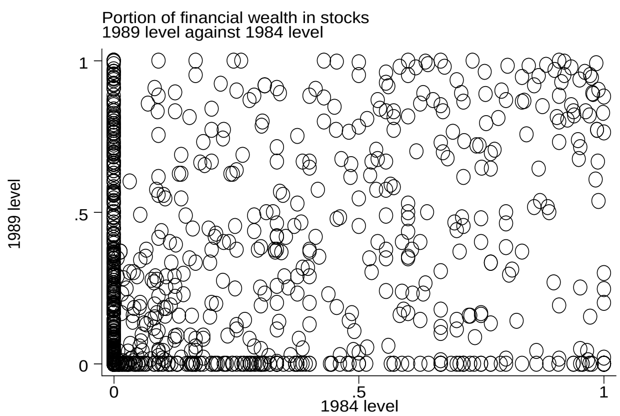

Fig.1b The fixed cost 'stock angle', 1984 and 1989 levels

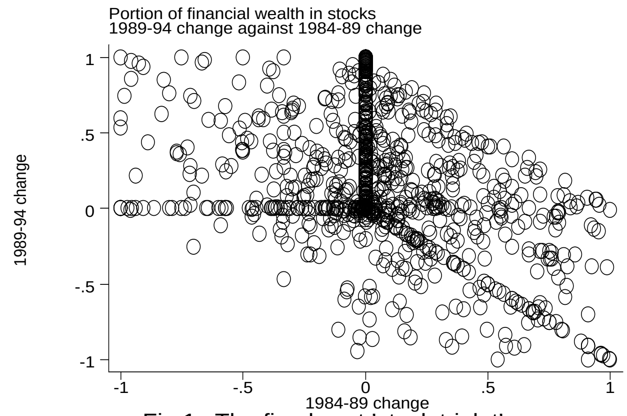

Fig.1c The fixed cost 'stock triplet'

Figure 1:

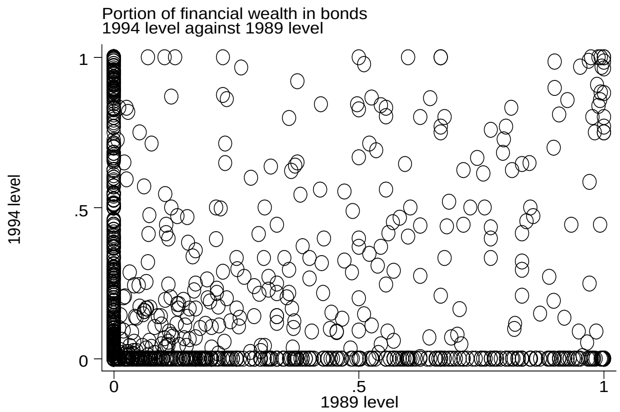

Fig.2a The fixed cost 'bond angle', 1989 and 1994 levels

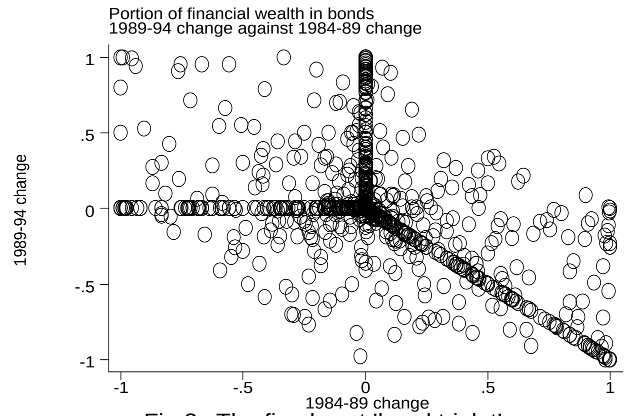

Fig.2c The fixed cost 'bond triplet'

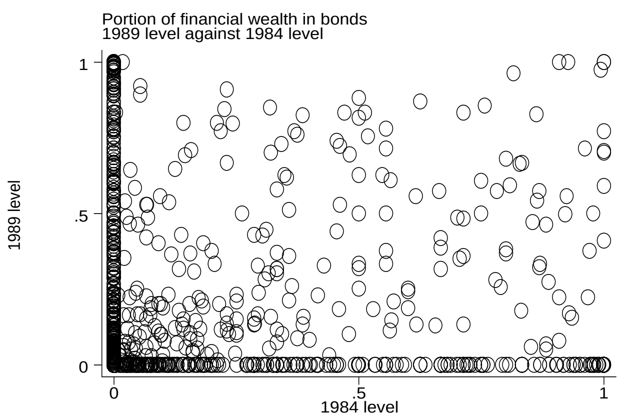

Fig.2b The fixed cost 'bond angle', 1984 and 1989 levels

Figure 2:

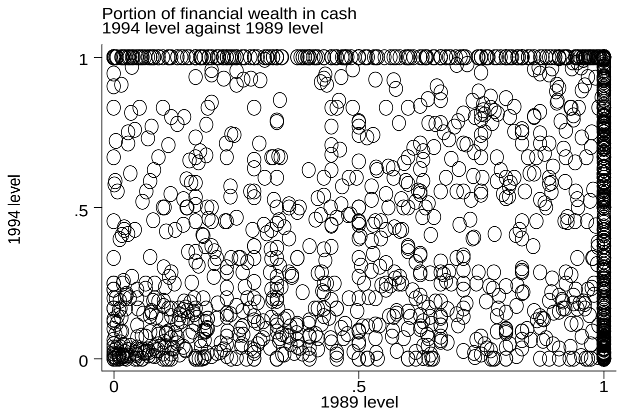

Fig.3a The fixed cost 'cash angle', 1989 and 1994 levels

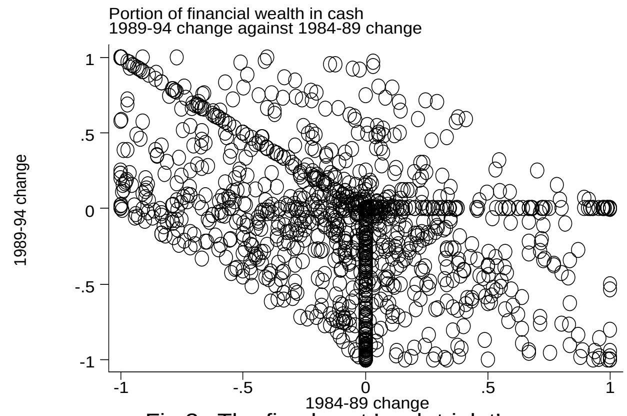

Fig.3c The fixed cost 'cash triplet'

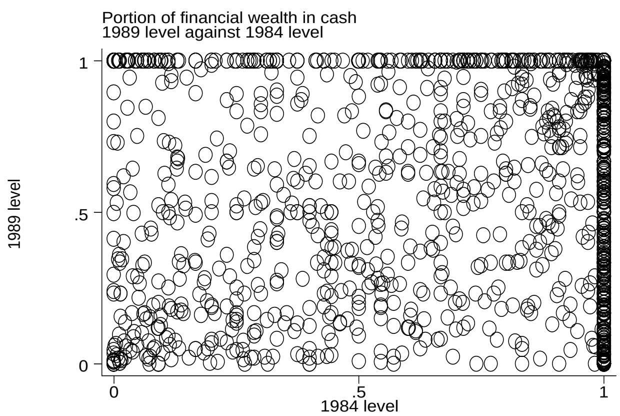

Fig.3b The fixed cost 'cash angle', 1984 and 1989 levels

Figure 3:

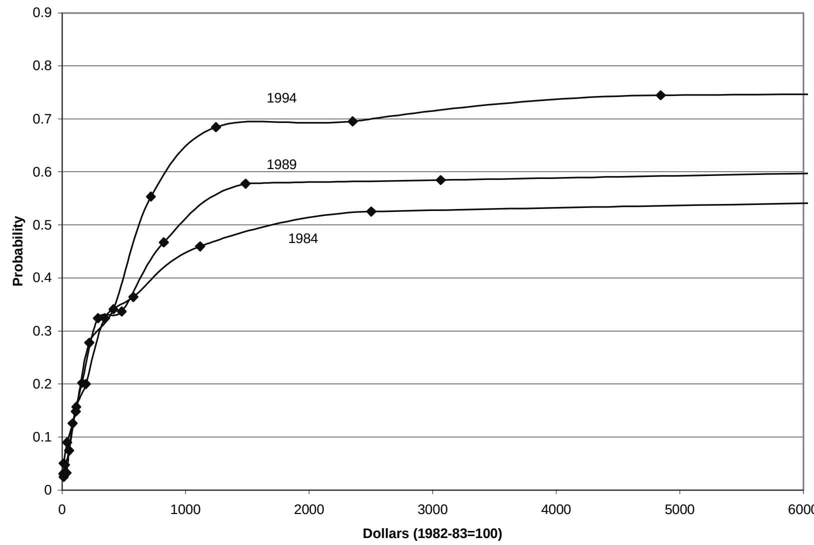

Figure 4: Estimated CDF of per period cost, $\alpha = 1$

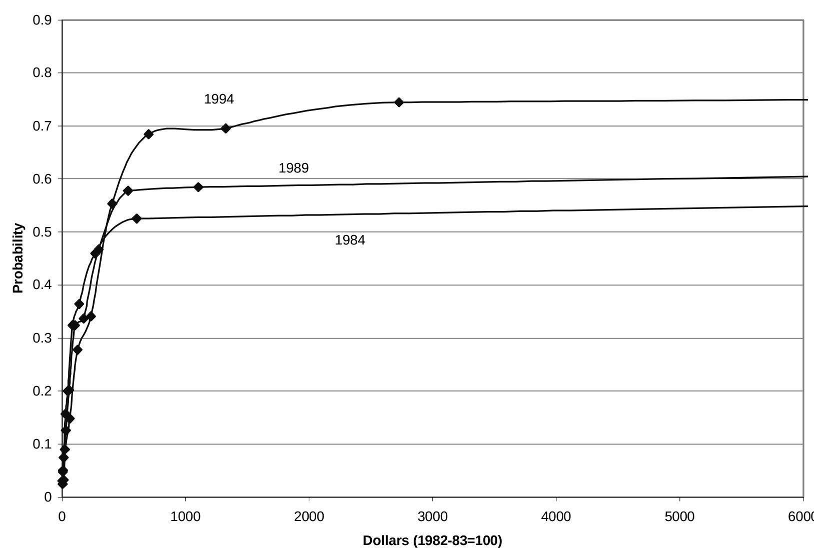

Figure 5: Estimated CDF of per period cost, α =period average