CSEF

Centre for Studies in Economics and Finance

WORKING PAPER NO. 156

The Foregone Gains of Incomplete Portfolios

Monica Pajella

February 2006

University of Naples Federico II

University of Salerno

Bocconi University, Milan

CSEF - Centre for Studies in Economics and Finance – UNIVERSITY OF SALERNO
84084 FISCIANO (SA) - ITALY
Tel. +39 089 96 3167/3168 - Fax +39 089 96 3167 – e-mail: csef@unisa.it

WORKING PAPER NO. 156

The Foregone Gains of Incomplete Portfolios♠

Monica Paiella♦

Abstract

This paper estimates a lower bound to the foregone gains of incomplete portfolios, which are in turn a lower bound to the (unobserved) entry costs that could rationalize non-participation to financial markets. My estimates provide a heuristic test for the cost-based explanation of limited financial market participation: high estimates would imply implausibly high participation costs. Using the CEX and assuming isoelastic utility and a relative risk aversion of 3 or less, for the stock market I estimate an *average* lower bound ranging between 0.7 and 3.3 percent of consumption. Since annual total (observable plus unobservable) participation costs are likely to exceed these bounds, the cost-based explanation is not rejected by this test.

JEL Classification: G11, D12, E21

Keywords: intertemporal consumption model, financial market participation, household portfolio allocation, non-proportional cost of participation, near-rationality

♠ This paper is a substantially revised and improved version of my working paper "Limited Financial Market Participation: A Transaction Cost-based Explanation" (Institute for Fiscal Studies WP01/06). Thanks are due to Orazio Attanasio, Richard Blundell, John Cochrane, Janice Eberly, Dario Focarelli, John Heaton, Erzo Luttmer, Ellen McGrattan, Luigi Pistaferri, Andrea Tiseno, Annette Vissing-Jørgensen for suggestions and discussion and to participants to the seminars at the University College London, Institute for Fiscal Studies, Bank of Italy, Ente Einaudi and Università di Salerno, and to the NBER Summer Workshop on the Aggregate Implications of Microeconomic Behavior. The views expressed are those of the author and do not necessarily reflect those of the Bank of Italy. The author is also the only one responsible for any errors

♦ Bank of Italy, Research Department

Address for correspondence: Bank of Italy, Economic Research Department, Via Nazionale 91, Roma 00184, Italy, tel. +39.06.4792.2595, fax. +39.06.4792.3723. E-mail: monica.paiella@bancaditalia.it.

Contents

A large number of studies have suggested that observed asset returns are inconsistent with consumption choices as predicted by the theory of the intertemporal allocation of consumption. Recent studies have explored the possibility that limited participation in financial markets might explain the disparity between theoretical predictions and empirical evidence. More precisely, since the first order conditions of asset pricing models hold with equality only for those households who own complete portfolios, the models should be tested for this subset of households and not for the whole population. As a consequence, since in practice relatively few households hold shares directly, the use of aggregate consumption data in evaluating asset pricing models could be very misleading.[1] These are the points stressed by Mankiw and Zeldes (1991), Attanasio, Banks and Tanner (2002), Paiella (2004) and Vissing-Jørgensen (2002), among others, who propose limited financial market participation as a unified framework for rationalizing the empirical rejections of the Consumption-based capital asset pricing model. These papers show that accounting for portfolio heterogeneity helps reconciling the predictions of the theory with empirical evidence. These studies take limited participation as given and make no attempt of rationalizing it. However, limited participation is in itself a puzzle for the intertemporal consumption model. Merton (1969) and Samuelson (1969) have illustrated how such behavior is inconsistent with the maximization of expected lifetime utility, which predicts that expected-utility maximizers should always be willing to invest an arbitrarily small amount in all assets with positive expected return, including the risky ones, unless there are non-linearities in the budget constraint. Yet, among households surveyed in the US Consumer Expenditure Survey over the period 1982-95, more than two thirds do not hold either stocks or bonds.[2] This behavior is usually reconciled with the intertemporal consumption model by invoking non-proportional costs of financial market participation (explicit and non-explicit). How plausible this explanation is, depends crucially on how large the theory predicts that these costs should be, relative to their actual size in reality which is to a large extent unobservable (for example, the value of time spent gathering information is not observable).

This paper focuses on the issue of limited participation to financial markets and proposes a heuristic test of the plausibility of the "participation cost explanation". Using the empirical implications of the consumption model, I estimate a lower bound to the foregone gains of holding an incomplete portfolio: by construction, the true foregone gains are higher than my estimated bound. The true foregone gains are in turn a lower bound to the level of costs that would rationalize non-participation. Hence, through this double inequality I identify a lower bound to the costs that could reconcile incomplete portfolios with the intertemporal consumption model: any such cost must be higher than my estimated bound. The estimated bound provides a heuristic test of the plausibility of the "cost-of-participation story". If the estimated bound were very high, the cost-of-participation explanation would be implausible because costs would be unlikely to exceed such level in reality. My bound estimate varies with risk aversion and is increasing in consumption and wealth. Given a relative risk aversion equal to 3 or less, which is the range of values considered plausible by the literature on the equity premium puzzle,[3] and averaging across *all* riskless asset holders, I find that the lower bound to the foregone gains for non-investing in risky assets ranges between 0.7 and 3.3 percent of expenditure on non-durable goods and services. However, for the wealthiest third of riskless asset holders, the foregone gains could be as high as 6.7 percent. The bound to the foregone gains for non-investing in riskless assets ranges between 0.3 and 1.1 percent, on average.

Overall, the bounds are quite precisely estimated. Nevertheless, since this study uses microdata for consumption, measurement error may be a serious problem. As a matter of fact measurement error tends to make household consumption appear less smooth than it actually is. This is a potential source of upward bias for my estimates of the gains. In fact, it may lead to an overstatement of the benefits from investing in the "overlooked" security, as it can be used to make consumption smoother. A Monte Carlo analysis of its implications for the estimation of the benefits of investing in risky and riskless assets suggests that measurement error in consumption is indeed a source of bias. However, on the whole, the bias appears to be relatively small, especially when focusing on investments in risky assets and for low levels of risk aversion. In fact, as expected, the bias is increasing in risk aversion, as the more risk averse value consumption smoothing relatively more. More importantly, measurement error biases the bound estimates upwards. Hence, the true bounds can be expected to be lower than the ones that I find.[4]

As to the true costs of participation to financial markets, which include both a monetary and non-monetary components, according to the US National Income and Product Accounts, monetary charges for brokerage and trust services and investment counselling in 1996 amounted to 2.9 percent of consumption (3.7 percent in 2000).[5] Non-monetary charges refer to the total value of time spent acquiring and processing information and managing the optimal portfolio. Given an average consumption-income ratio of 0.5, spending just one hour per week in caring for one's investment in the stock market would add up to a cost of 5 percent of consumption.[6] These figures suggest that the true costs of participating to the stock market are likely to be larger than my estimates of the average bound, supporting the claim that the heuristic test does not reject the "participation cost hypothesis", even at a reasonably high degree of confidence. However, the heterogeneity of the gains and the fact that the wealthiest non-participants appear to forego substantially larger returns imply that some caution is due when applying the participation cost hypothesis to explain non-participation among those with substantially large amounts of financial wealth.

The estimation of the lower bound to the annual foregone gains is based on the necessary conditions for the optimality of observed behavior of non-participants. The methodology relies on the (empirical) distinction between the consumption paths of those households holding a well-diversified portfolio and the consumption paths of those holding incomplete portfolios. Two different types of "incompleteness" are considered: portfolios of riskless assets, but without risky assets, and portfolios without risky and riskless assets.[7] This approach builds on a paper on adjustment costs and asset pricing by Luttmer. Luttmer (1999) focuses on the losses for leaving unexploited some trading opportunities and proposes a lower bound on the level of fixed transaction costs reconciling per-capita expenditure and asset returns. Hence, the foregone gains that Luttmer identifies bound from below the cost of trading that would justify not taking advantage of temporary changes in returns not matched by changes in the riskiness of assets. Instead, I use individual level data and, by distinguishing between holders and non-holders of risky assets, I focus on the loss for missing out on the equity premium by those who only use sub-optimal portfolios to smooth consumption over time. By distinguishing between holders and non-holders of riskless assets I also consider the loss for missing out on the riskless rate by those who use other means (such as durables, currency or something else) to smooth consumption over time. Luttmer's frictions are the costs that a representative agent must pay to trade and modify her consumption path. My frictions are the costs that individual agents must pay in order to participate to financial markets. Such participation costs can be thought of as reflecting the costs of information and transaction that would induce households not to invest in some securities. Alternatively, they can be thought of as the costs of following "near-rational" decision rules. The idea is that, in practice, households do not literally maximize their utility, but follow heuristic decision processes that economists model as solutions to a maximization problem. Households' decisions deviate from the solutions to this problem, but with limited costs in terms of utility (for a thorough discussion of this view, see Simon (1978)).

Two related papers are Cochrane (1989) and Vissing-Jørgensen (2003). The former analyses the sensitivity of the tests of the intertemporal allocation of consumption to near-rational alternatives and estimates somewhat simpler utility costs as a measure of "economic standard errors" for different predictions of a model. The latter provides evidence on the distribution of the per-period participation costs in the cross-section by looking at the dollar gains from moving a fraction of household's financial wealth into the stock market, assuming unchanged current period consumption. Vissing-Jørgensen's results are fully consistent with the evidence that I present in the paper, with her estimates of the per-period participation cost being higher than my bounds and the cost increasing in household wealth.

The rest of the paper is organized as follows. Section 1 presents the framework for bounding the gains that an agent holding an incomplete portfolio of assets forgoes within the type of environment specified by the model of intertemporal choice. The estimation procedure *vis à vis* the data available are examined in Section 2. Section 3 presents estimates of the gains that riskless asset holders forego for non-investing in risky assets. Section 4 considers the forgone gains of those who do not hold financial assets at all. Section 5 concludes.

1. Measuring the Foregone Gains

The framework developed in this section responds to the desire of rationalizing the choice of non-investing in some of the assets available, despite the apparent sub-optimality of this behavior. The analysis is based on the idea that consumers holding incomplete portfolios could improve their consumption path by investing in the assets they do not own, but this would involve information gathering, decision-making, brokerage fees and other costs, creating a disincentive to complete portfolio diversification. Such frictions end up preventing consumers from exploiting the equity premium paid by the risky assets. In some cases, they even prevent them from smoothing consumption by exploiting the riskless rate paid by the riskless asset. The paper tests this hypothesis by bounding the costs that would rationalize incomplete portfolios under standard assumptions regarding preferences. The bound is determined by comparing the utility associated to the choice of non-investing in some asset to the utility associated to the investment, using a revealed preference argument as follows.

Consider an environment where households have rational expectations, inter-temporally, additively separable preferences over consumption, a strictly increasing and concave per-period utility function, $U(.)$, and a positive subjective discount rate, β. Households have access to several means to substitute consumption over time. In particular, they can accumulate financial securities, currency and/or real assets. Some of these substitution opportunities are easier to use than others and, on the basis of portfolio composition, it is possible to distinguish among three types of households: those who hold both risky and riskless financial assets (type 1); those who hold only riskless assets (type 2); and those who hold neither (type 3). Let $\{c^h\}_t$, $t = 1, 2, ...$ be household *h* observable sequence of consumption choices. These choices result from some complicated and unobservable set of decisions involving labour supply, saving and portfolio allocation. Since households choose optimally, conditional on the information available, and, at time *t*, they could have chosen any other feasible sequence of consumption bundles, their time *t* expected utility gain from deviating from $\{c^h\}_t$ must be non-positive.

Let's focus on those who have chosen not to invest in some of the financial assets available. Consider a one-period asset with return R_{t+1} between period *t* and *t*+1. Suppose that, at time *t*, household *h* chooses not to invest in this asset. It invests its wealth $W_{h,t}$ (if any) in a portfolio yielding R^f_{t+1}, consumes $c_{h,t}$ and expects to consume $c_{h,t+1}$ at time *t*+1. Suppose that, instead, at time *t* this agent could have paid a cost of δ^* units of consumption, invested in the asset and consumed ($\tilde{c}_{h,t}$, $\tilde{c}_{h,t+1}$) instead of ($c_{h,t}$, $c_{h,t+1}$), with ($\tilde{c}_{h,t}$, $\tilde{c}_{h,t+1}$) defined as:

$$\tilde{c}_{h,t} = c_{h,t} - x^c_{h,t} c_{h,t} - \delta^* c_{h,t} , \tag{1}$$

$$\tilde{c}_{h,t+1} = c_{h,t+1} + x^c_{h,t} c_{h,t} R_{t+1} + x^w_{h,t} W_{h,t} (R_{t+1} - R^f_{t+1}) , \tag{2}$$

where $x^c_{h,t}$ and $x^w_{h,t}$ denote the fraction of time *t* expenditure and wealth that the household invests in the asset after paying the cost. Let $E_t\{\ \}$ denote the expectation conditional on the information available at time *t*, when deciding whether or not to pay the cost and use financial markets to adjust consumption. Optimality of observed choices ($c_{h,t}$, $c_{h,t+1}$) implies that:

$$E_t \left\{ U(\tilde{c}_{h,t}) + \beta U(\tilde{c}_{h,t+1}) \right\} \le E_t \left\{ U(c_{h,t}) + \beta U(c_{h,t+1}) \right\}. \tag{3}$$

Equation (3) says that, net of the cost δ^*, the expected utility gain from perturbing the observed consumption path is non-positive, hence not worth. Inequalities like (3) must hold for any *t*.

For ease of notation, let's write the foregone utility gain of household *h*, for the misallocation over the periods *t* and *t*+1, as:

$$v_{h,t+1}\left(x^c_{h,t}, x^w_{h,t}, \delta^*\right) = U(\tilde{c}_{h,t}) + \beta U(\tilde{c}_{h,t+1}) - U(c_{h,t}) - \beta U(c_{h,t+1}). \tag{4}$$

Then, (3) can be re-written as an inequality regarding a function of $x^c_{h,t}$, $x^w_{h,t}$ and δ^*:

$$E_t \left\{ v_{h,t+1}\left(x^c_{h,t}, x^w_{h,t}, \delta^*\right) \right\} \le 0. \tag{5}$$

If long individual consumption series were available, one could proceed as Luttmer (1999) and study (5) by taking unconditional expectations. Instead, the data used allow to determine only one observation on foregone gains per non-participating household. To deal with this complication, let H_t denote the set of agents, observed at time t, who do not invest in the asset considered, based on the information available at t. Then, summing the individual $v_{h,t+1}\left(x^c_{h,t}, x^w_{h,t}, \delta^*\right)$ over H_t, and taking unconditional expectations yields:

$$E\left\{\sum_{h\in H_t} v_{h,t+1}\left(x^c_{h,t}, x^w_{h,t}, \delta^*\right)\right\} \leq 0. \tag{6}$$

This inequality holds for any $x^c_{h,t}$ and $x^w_{h,t}$. Therefore, varying $x^c_{h,t}$ and $x^w_{h,t}$, one obtains:

$$\sup_{x^c_{h,t}, x^w_{h,t}} E\left\{\sum_{h\in H_t} v_{h,t+1}\left(x^c_{h,t}, x^w_{h,t}, \delta^*\right)\right\} \leq 0. \tag{7}$$

The left-hand-side of (7) is non-negative at $\delta^* = 0$, and continuous and decreasing in δ^* – as $U(.)$ is continuous and increasing. Hence, we can construct a lower bound d for δ^* by solving:

$$\sup_{x^c_{h,t}, x^w_{h,t}} E\left\{\sum_{h\in H_t} v_{h,t+1}\left(x^c_{h,t}, x^w_{h,t}, d\right)\right\} = 0, \tag{8}$$

which is such that (7) is satisfied for any $\delta^* \geq d$.

The parameter d is the Hicks compensating variation for not investing in an asset yielding R_{t+1}. This is a *lower* bound to the foregone gains for holding an incomplete portfolio, which in turn are a lower bound to the costs that would rationalize non-participation. The "true" foregone gains for holding an incomplete portfolio are just a *lower* bound to the participation costs, because the (unobservable) costs δ^* may be so large that households are never close to deviating from their actual choices. In this instance, by construction, a level of gains that is much smaller than δ^* will suffice to rationalize observed choices. Further, and more importantly, by construction, the parameter d is only a lower bound to the foregone gains of incomplete portfolios: the expected utility gains of deviating from observed portfolio choices may be higher than those captured by equation (4) for at least two reasons. The framework behind equation (4) measures the expected gains of using an extra instrument to adjust consumption over two periods. Thus, first of all, if the conditioning information set of the agent is larger than that of the econometrician, the agent may be actually able to obtain a higher utility gain than the econometrician can estimate. Secondly, the type of behavior implied by equation (2) is not optimal as it implies consuming at time t+1 all the return on the investment. Nevertheless, this set up leaves households' consumption plans unchanged at all other dates and allows to appraise the gains that households forego for non-investing for one period by focusing just on their consumption at two adjacent dates. Optimal behavior would allow for the gains from the investment to be spread over the entire lifetime horizon of the utility maximizing agent. My setup approximates the utility from such stream of gains with the utility from the gains over the two periods when the investment takes place, assuming that the entire return from the investment is consumed at t+1.

Overall d provides the basis for a heuristic test of the cost of participation hypothesis: for the latter to be a plausible explanation of incomplete portfolios, any reasonable cost of participation must be higher than my estimated bound. Although this is not the most powerful test, it is indeed the most reliable. A more powerful test would compare the costs with the true foregone gain –not just with a lower bound, as done here. However, the estimation of the true foregone gain would require a much larger amount of information and/or assumptions.

2. Estimation Strategy

2.1. Empirical specification and estimation procedure

The trading rules in case of participation are assumed to be linear in consumption and wealth. The trading rule for consumption is specified as: $x_{h,t}^{c} = x_{t}^{c}(z_{t}; a_{c}) = a_{c}' z_{t}$. This set up allows capturing in the estimate the predictability of the components of asset returns that are correlated with consumption growth and the set of forecasting variables z_t, which help selecting the most profitable level of investment in case of participation. Short-sale constraints can be taken into account by imposing $a_{c}' z_{t} \geq 0$ for all t. The trading rule for wealth is specified as: $x_{h,t}^{w} = x^{w}(a_{w}) = a_{w}$. Using forecasting variable in the trading rule for wealth does not affect the estimation of the bound, but would significantly increase the computational burden.

Assume that the supremum of the foregone gain function in (6) is attained and let (α_c, α_w) denote the optimal values of (a_c, a_w). The estimation of the lower bound to the foregone gains of incomplete portfolios relies on the following set of first-order conditions:

$$E\left\{ \sum_{h \in H_t} D_j v_{h,t+1}\left(x_t^c(z_t; \alpha_c), x^w(\alpha_w), d \right) \right\} = 0, \quad j = 1, 2; \tag{9}$$

$$E\left\{ \sum_{h \in H_t} v_{h,t+1}\left(x_t^c(z_t; \alpha_c), x^w(\alpha_w), d \right) \right\} = 0, \tag{10}$$

where D_j denotes the derivative with respect to the j^{th} element of $v_{h,t+1}\left(x_t^c(.), x^w(.), d \right)$. Equations (9)[8] determine the optimal investment in case of participation given the cost. Since, in practice, the actual cost δ^* is not observed, nor estimated and only a lower bound to the cost is identified, the optimal portfolio is determined as a function of a cost equal to its estimated bound, *d*, which is consistent with the rest of the analysis. Equation (10) determines the lower bound *d* to the participation cost δ^*, given the optimal investment.

To determine the fixed-cost bound *d*, one can use a method of moment estimator based on the sample analogues of (9) and (10). In practice, to simplify the estimation, I determine α_w by grid search, with the grid going from 0 (no wealth is moved into the risky asset) to 1 (all wealth is moved), at 0.05 intervals. The optimal wealth share is the one that maximizes the gains in case of participation. Given a_w, I replace the expectations in (9) and (10) by sample averages and use for the estimation the following conditions:

$$\frac{1}{T-1}\sum_{t=1}^{T-1}\left\{\sum_{h\in H_t} D_1 v_{h,t+1}\left(x_t^c(z_t;\alpha_c), x^w(a_w), d\right)\right\} = 0; \tag{11}$$

$$\frac{1}{T-1}\sum_{t=1}^{T-1}\left\{\sum_{h\in H_t} v_{h,t+1}\left(x_t^c(z_t;\alpha_c), x^w(a_w), d\right)\right\} = 0. \tag{12}$$

The estimator of *d* is consistent in *T* and only if the investment rule, $x_t^c(z_t;a_c)$, as a function of the parameters is well behaved and if $\left\{\sum_{h\in H_t} D_1 v_{h,t+1}\left(x_t^c(z_t;\alpha_c), x^w(a_w), d\right)\right\}$ and $\left\{\sum_{h\in H_t} v_{h,t+1}\left(x_t^c(z_t;\alpha_c), x^w(a_w), d\right)\right\}$ are time stationary and have finite mean, so that some law of large numbers can be applied.

For the estimation, I assume that household utility exhibits constant relative risk aversion, i.e.:

$$U(c_{h,t}) = \frac{c_{h,t}^{\;1-\gamma} - 1}{1-\gamma}; \tag{15}$$

where $c_{h,t}$ is household consumption of non-durable goods and services in period *t*. Notice that neither the coefficient of relative risk aversion γ, characterizing household isoelastic preferences, nor the subjective discount rate, β, can be identified within the model. Hence, I assign them a range of values and verify how sensitive the estimates of the optimal portfolio and of the foregone gains are to such parameters.

In practice, on the basis of the data available, I can identify the bounds to the costs that would rationalize two types of incomplete portfolios. First, I consider the foregone gains for not holding risky assets, but only riskless ones (type 2 households), and bound the cost of non-investing in a well-diversified portfolio of risky assets. Second, I look at the foregone gains of type 3 households. Using the information on these households, I can determine a bound to two different costs: the cost that would justify the choice of non-investing even just in riskless assets and the cost that would justify the choice of non-investing in an optimally determined portfolio of risky and riskless assets.[9]

2.2. Data

The data used are taken from the US Consumer Expenditure Survey (CEX), which is a representative sample of the US population, run on an ongoing basis by the US Bureau of Labor Statistics. The CEX has a rotating panel dimension, with each consumer unit being interviewed every three months over a twelve month period. As households complete their participation, new ones are introduced into the panel on a regular basis and, as a whole, about 4500 households are interviewed each quarter, more or less evenly spread over the three months.

At the time of the last interview, households provide information on their asset holdings at that date and on the “dollar difference” with respect to the amounts held twelve months earlier. The asset categories in the CEX are: 1. checking, brokerage and other accounts; 2. saving accounts; 3. US savings bonds; 4. stocks, bonds, mutual funds and other securities. As a measure of risky asset holdings, I take the amounts held in “stocks, bonds, mutual funds and other securities” and “US saving bonds”; as a measure of riskless asset holdings, I take the amounts held in checking and saving accounts. In order to determine the asset holding status at t, for each asset category, I subtract from the stocks held at the time of the last interview (t+1) the amount of savings (the dollar change) carried out over the previous twelve months. Table 1 reports the sample composition in each of the years considered on the ground of household asset portfolios. On average, 31 percent of the sample holds positive amounts of both risky and riskless assets; 50 percent holds only riskless assets; 19 percent does not hold either asset. In the sample used, no household holds only risky assets. The evidence reported in the Table suggests that the share of households owning stocks and bonds has increased substantially over the years covered by the survey, which is consistent with the evidence from other data sets.

Each quarterly interview collects detailed information on expenditure for each of the three months preceding the one when the interview takes place. However, since the information on asset holdings is annual and I consider household portfolios at the time when they enter the survey, for each household I can define only one observation on the expected utility gain. Hence, I use only two observations on consumption, with $c_{h,t}$ and $c_{h,t+1}$ denoting spending in the months preceding the first and the last interview, respectively.[10] The consumption measure that I use is deseasonalized, real monthly, per-adult equivalent expenditure[11] on non-durable goods and services. For consistency, financial wealth is also rescaled to real, per-adult equivalent terms and is divided by 12 for comparability with the (monthly) consumption measure.

The data used for the analysis cover the period 1982-1995.[12] Around 1985-86, the sample design and the household identification numbers were changed and after the first quarter of 1986 no track is kept of the households who had entered the survey in 1985. As a consequence of this and of the fact that the information on financial asset holdings is collected during the last interview, those households who have their first interview in the third and fourth quarter of 1985 had to be excluded from the sample. Thus, the sample used consists of households who have their first interview between 1982:1 and 1985:6 and between 1986:1 and 1995:1 and t runs for a total of 150 periods.

From the initial sample of households, I exclude those with incomplete income responses and those whose financial supplement contains invalid blanks in either the stocks of assets held or in the dollar changes occurred with respect to the previous year. In addition, I exclude those living in rural areas or in university housing, those whose head was less than twenty-five or older than sixty-five and those who do not participate to all interviews (about 33% of the initial sample). I then select out the top 0.1 percent of the income distribution and the bottom 1.7 percent, which corresponds to about 500 households whose total after-tax annual income is below \$3,500 and who are likely to consume all their income and have no resources to invest in financial markets. More importantly, these households are likely to be financially constrained and for the liquidity constrained the standard conditions for optimality of behavior, upon which my analysis builds, do not hold. I exclude also those households with average monthly per-adult equivalent consumption lower than \$250 (about 1,000 households corresponding to 3.6% of the sample) and those who report a change in per-adult equivalent consumption, $\Delta c_{h,t}$, greater than \$1,750 in absolute value (about 500 households). Finally, I drop those households holdings only risky assets (less than 0.4% of the sample). Overall, the sample used consists of 23,970 households.

Table 2 reports some descriptive statistics for the sample as a whole and for the three types of households. Type 1 households, who hold both risky and riskless assets, are more likely to be headed by a man, the household head is more educated than the average, slightly older and more often married. Their after-tax income and consumption are also relatively higher. Those who hold neither risky nor riskless assets (type 3) tend to be the least educated and to have the lowest income and consumption and in 41% of the cases are headed by a woman.

Real annual asset returns are summarized in Table 3. The risky return corresponds to the total return (capital gains plus dividends) on the S&P500 Composite Share Index. The riskless return coincides with the return on 3-month US Treasury bills. The mean equity premium that those who do not invest in the risky asset forego is about nine percentage points.

3. The Foregone Gains for Non-Investing in Risky Assets

Tables 4 through 7 present estimates of the gains that riskless asset holders forego when they choose not to invest in a well-diversified risky portfolio whose return mimics that on the S&P500 Composite Index. The share of consumption to be invested in the asset market after paying the cost, $x_t^c(z_t; a_c)$, is assumed to be a function of a vector of instruments (z_t), that have been shown to be useful in predicting asset returns (see, for example, Kleim and Stambaugh (1986) and Fama and French (1989)). The instruments include the returns on the S&P500 CI and on 3-month Treasury bills, the term spread, and the price-earnings ratio, plus a constant. Variables are lagged one period and refer to the time interval denoted as '*t*-1–to–*t*'. The tables have the following basic structure. Each column is computed assuming isoelastic preferences for different levels of risk aversion. The rate of discount, β, is set equal to 0.98.[13] Tables 4 and 5 are obtained by averaging the foregone gains over all riskless asset holders. In Tables 6 and 7, households are sorted based on their initial consumption and wealth, which allows to appraise the extent of the differences in the gains related to these characteristics.

Table 4 is based on the assumption that riskless asset holders do not move any of their riskless wealth into the risky asset ($x^w(a_w) = 0$). Panel (a) considers the case where households do not attempt to time the market when determining their optimal investment, which implies a constant trading rule for consumption. When risk aversion is equal to 1, the optimal consumption share to invest corresponds to around 7 percent of consumption. Consistent with the literature on portfolio choice, it is decreasing in risk aversion and drops to 4 percent when γ is 3 and to under 3 percent when γ is 10. As to the foregone gains, for a riskless asset holder with a relative risk aversion of 1, and investing optimally, they amount to at least 0.6 percent of her expenditure. As γ increases, the point estimate of d rises, although the differences are hardly statistically significant. For values of γ equal to 8 or higher, it becomes statistically insignificant. The tendency to increase can be explained on two grounds. First of all, some of the increase is likely to be due to the estimation bias induced by mis-measured consumption. As discussed earlier, such bias is increasing in risk aversion. Second, the benefits of investing in a well-diversified risky portfolio are related not just to the exploitation of the equity premium, which benefits the most the least risk averse, but also to the availability of an effective means to smooth consumption over time, although at a rather high risk. When the benefits come from an improved intertemporal reallocation of expenditure, the foregone gains can be expected to increase in risk aversion, because the risk averse value consumption smoothing relatively more. The fact that the estimated d exhibits a tendency to increase in γ suggests that the benefits of smoothing consumption with a high return, well-diversified risky portfolio may be important and large enough to offset the loss due to the higher volatility of $c_{h,t+1}$.

Panel (b) of Table 4 considers the case where non-participants try to predict excess returns when determining their optimal investment in the risky asset. Despite their strong significance in regressions predicting the equity premium, with the exception of the price-earnings ratio, which is always statistically significant, different variables seem to matter at different levels of risk aversion and, in each instance, some of the instruments turn out to be scarcely significant at the standard levels. Overall, timing the market appears to increase significantly the gains from participation for the highly risk averse and makes the corresponding bound estimates statistically significant. However, the precision of the estimates for high values of γ is low, as the standard errors are quite high. The bound estimates corresponding to low or intermediate levels of risk aversion are to a large extent unaffected.

The results in Table 5 are based on the assumption that households may find it optimal to move some of their wealth from the riskfree into the risky asset, once they pay the fixed cost. As before, the overall investment in risky assets is estimated to be decreasing in risk aversion and participation is profitable even for the most risk averse only if households time the market when determining the optimal trading rule for consumption (panel (b) of the table). The optimal consumption share to invest is decreasing in γ and goes from 5 to 2.5 percent. The optimal fraction of wealth to reallocate turns out to be very high and ranges from 100 percent for the least risk averse to 10 percent. This result is consistent with the evidence of

Heaton and Lucas (1997, 2000) who simulate household portfolio allocation and find that for moderate-to-low levels of risk aversion agents hold only stocks almost all of the time. When households invest in the stock market also some of their wealth, the gains from participation related to the exploitation of the equity premium increase substantially, but also the volatility of t+1 consumption. Overall, with respect to the case where $x^w(a_w) = 0$, the estimated foregone gain is significantly higher, but only for the little risk averse. When $\gamma = 1$, d is 5 times higher than when $x^w(a_w)$ is set to 0 and equal to 3.3 percent; for $\gamma = 3$, it is twice as large and equal to 1.8 percent, but for $\gamma > 6$, the difference is negligible. The bound estimates reported in this table should be taken with some caution. In fact, it is unlikely that riskless asset holders would *actually* invest such larger shares of their financial wealth in the risky asset. Heaton and Lucas (2000) show that, when background risk (which is ignored here) in the form of labor and proprietary income risk and real estate risk are properly accounted for in the simulation, the average portfolio share in stocks declines substantially, which in my framework would translate into somewhat lower gains. If I set $x^w(a_w) = 0.3$, which is the mean portfolio share in stocks in my sample, the estimated gain would drop to 1.7 percent when $\gamma = 1$, and to 1.5 when γ is 3.[14]

The gains that households forego when they choose not to invest in stocks can be expected to be higher the more the resources available for investment. The analysis carried out so far allows for some across household heterogeneity in the foregone gains by stating the bound as a percentage of consumption. Tables 6 and 7 investigate further the issue of gain heterogeneity by splitting non-shareholders into three groups based on the size of initial wealth and two groups based on the size of initial consumption. This allows to tighten the cost bound estimate by focusing on those whose gains are highest. Table 6 is obtained by setting $x^w(a_w) = 0$, whereas in Table 7 a_w is estimated (by grid search) and set optimally. Households are assumed to time the market when determining the trading rule for consumption using the same instruments listed earlier. Only average consumption shares are reported in the tables (trading rule coefficient estimates are available upon request).

When $x^w(a_w) = 0$ (Table 6), not surprisingly, the bound estimates do not vary across the wealth distribution, as the gains come from investing only a fraction of consumption. Instead, the bound is increasing in consumption: the higher consumption, the smaller the marginal reduction in time t utility associated to the investment (and to the payment of the cost) which increases the amount households will invest and the overall gain from participation. Overall, the estimated bound is statistically significant generally only if risk aversion is below 5, which might be due to the sample sizes, which are relatively small, together with the fact that the standard errors are increasing in γ. When statistically significant, values range between 0.5 and 1.2 percent of consumption. When households can move their wealth from the riskless into the risky asset, and $x^{w}(a_{w})$ is determined optimally (Table 7), important differences in the size of the foregone gains emerge across the groups considered, with the gains sharply increasing in investors' resources. Overall, for those in the bottom two thirds of the financial wealth distribution, the bound estimate is under 2 percent. For those in the top third of the distribution, the foregone gains are much higher. For γ=1, the bound estimate may be as high as 6.7 percent; for γ=3 it drops to around 3.5 percent.[15] Like before with Table 5, a word of caution is due because these results crucially depend on the fact that households move a substantial fraction of their riskless wealth into the risky asset and α_w does not account for any background risk households may be subject to. Business income and real estate risks may be particularly important especially for the wealthiest and may make them behave in a more risk averse way, which would result in smaller stock investments. Nevertheless, these results confirm that relatively higher costs are needed to explain the choices of the wealthiest non-participants, which is fully consistent with the evidence of Vissing-Jorgensen (2003).

4. The Foregone Gains for Non-Investing in Financial Assets

Tables 8 and 9 are based on a set of households who do not hold neither risky nor riskless assets and on the assumption that these households smooth consumption using other means, on which however no information is available. Table 8 focuses on the foregone gains for non-investing in an asset yielding the three-month Tbill rate of return. In panels (a) and (b) households can only *buy* the asset; in panel (c) short-selling, i.e. borrowing at the riskless rate, is allowed. In panel (a) households do not time the market; in panels (b) and (c), they do so using lagged returns on the riskless and risky assets. The results are very similar across the three panels and are robust to using different sets of instruments to predict returns. According to the evidence in the table, a household with no financial assets and a risk aversion of 1 could increase its utility by investing around 5 percent of its consumption for one period in the riskless asset. As before, as risk aversion increases, the utility maximizing investment decreases. For high values of γ, there is some evidence that optimal behavior may involve some short-selling of the riskless asset, which may be a symptom of some type of liquidity constraint. Hence, at high values of risk aversion, results should be appraised with care. Overall, the estimated foregone gains turn out to be quite small and, as expected, are smaller than those recorded for risky asset market non-participation. The gains are increasing in risk aversion, but for $\gamma > 5$ they are statistically negligible. For low-to-moderate levels of risk aversion, they range from 0.3 to 2.6 percent of expenditure on non-durable and services. Note that allowing for borrowing does not increase the gains in any significant way.

Table 9 considers the case where type 3 households are allowed to invest in both risky and riskless assets. In this instance, the empirical specification for consumption in case of participation becomes:

$$\widetilde{c}_{h,t} = c_{h,t} - x^c(a_c)c_{h,t} - \delta^*_{r,rf}c_{h,t}; \quad (16)$$

$$\widetilde{c}_{h,t+1} = c_{h,t+1} + x^c(a_c)c_{h,t}R_{t+1} + y(a_y)c_{h,t}(R_{t+1} - R^f_{t+1}), \quad (17)$$

where $\delta^*_{r,rf}$ denotes the cost that household *h* has to pay to participate to a market where it can trade risky and riskless assets. $x^c(a_c)$ denotes the fraction of time *t* expenditure the household is willing to give up and invest in financial assets after paying the cost. $y(a_y)$ determines the allocation between risky and riskless assets. The investments in risky and riskless assets are equal to $\left(x^c(a_c) + y(a_y)\right)c_{h,t}$ and $-y(a_y)c_{h,t}$, respectively. When $y(a_y) < 0$, both assets are bought in positive amounts.[16] When $y(a_y) = 0$, there is no investment in the riskless asset; when $y(a_y) > 0$, the household borrows at the riskless rate to invest in the risky asset. The results in panel (a) of Table (9) have been obtained by imposing no borrowing; in panel (b) and (c) borrowing is allowed. In all cases, $x^c(a_c) \geq 0$, which implies that households consume all the return on the investment at time *t*+1. Allowing the investment functions to depend on vector of instruments does not change the results in any significant way.

The results in panel (a) of Table (9) suggest that, if type 3 households were to participate to financial markets and could choose to invest in either risky or riskless assets or in both, for low-to-moderate levels of risk aversion, their utility maximizing portfolio would consist basically just of risky assets. As before, the consumption share that households find optimal to invest is decreasing in risk aversion and ranges from 7.5 to 2.4 percent. Notice that the overall investment would be higher than that in riskless assets, when the latter were the only asset available (Table 8), which may suggest that the substitution effect prevails on the income effect. The estimated foregone gains are increasing in γ and are also generally higher than those for non-holding just riskless assets, reported in Table 8, as a result of the higher return on the overlooked investment. Overall, they range between 0.7 and 2.7 percent and for $\gamma>5$ they are statistically insignificant.

The foregone gain estimates turn out to be much larger when households are allowed to take short positions in the riskless asset to invest in stocks, as panel (b) of Table 9 shows. A household with log utility would maximize its participation gains if it could invest in the risky asset a sum equal to almost three times its consumption and finance the investment at the riskless rate. In this instance, its foregone gains would amount to 13 percent of its consumption. However, such levels of borrowing are unlikely to occur in practice, for two reasons. First of all, background risk is likely to make households behave in a more risk averse way, which significantly reduces borrowing and the gains from participation. As risk aversion increases, the optimal short position decreases rapidly together with the size of the foregone gains. For $\gamma=3$, the optimal investment in the risky asset would imply borrowing a sum equal to 88 percent of consumption and the corresponding gain would be 4.2 percent. For $\gamma>8$, the gains are statistically negligible. Secondly, it seems unlikely for households to be able to borrow at the 3-month TBill rate to finance a stock investment. Over the period considered, the mean rate charged on credit card balances, which can be taken as an indicator of the rates charged on non-collateralized loans, has amounted to 2.9 times the rate on 3-month TBills.[17] Panel (c) considers the case where households can finance their risky asset investment by borrowing at a rate which is set equal to twice the riskfree. In this case, for $\gamma=1$, the optimal investment in the risky asset would imply borrowing a sum approximately equal to household expenditure and the corresponding gain would be 1.8 percent. When $\gamma\geq5$, households would not find it optimal to borrow to invest in the risky asset.

5. Concluding Remarks

Non-participants' portfolio choices appear to depart significantly from utility maximizing behavior. The extent of the gains that households forego for holding incomplete portfolios depends on their degree of risk aversion, on the type of investment being considered, on whether they own other financial assets and, if they do, on whether they are willing or able to modify their wealth composition. For riskless asset holders, given a relative risk aversion of 3 or less, which is the range deemed plausible by the literature on the equity premium puzzle, I find that the *average* foregone gain for overlooking the possibility of investing in risky assets for one year ranges between 0.7 and 3.3 percent of household spending on non-durable goods and services. The foregone gain is increasing in household financial wealth and crucially depends on the optimal portfolio share in stocks. For the wealthiest third of the population, the foregone gains could be as high as 6.7 percent.

Using the information on those who hold neither risky nor riskless assets, it is possible to have a sense of the foregone gains for non-investing in riskless assets and in an optimally determined portfolio of risky and riskless assets. The foregone gain for non-investing in riskless assets is increasing in risk aversion and ranges between 0.3, if risk aversion is 1, and 2.6 percent of their consumption, if risk aversion is 5 (for higher levels of γ, it is statistically insignificant). When allowed to invest in both risky and riskless assets, households strongly prefer the former. If they were allowed to borrow at the riskless rate and invest in the risky asset and risk aversion sufficiently low, they would borrow up to three times their consumption and their foregone gains could be as high as 13 percent. However, such levels of borrowing are unlikely to occur in practice. On the one hand, background risk is likely to make households behave in a more risk averse way, which significantly reduces borrowing and the gains from participation: when risk aversion is 3, the foregone gain estimate drops to 4.2 percent. On the other, and more importantly, it seems unlikely for households to be able to borrow at the riskless rate. With borrowing at a rate twice as high as that on 3-month Tbills, the foregone gains would drop to below 2 percent.

These estimates are a structural implication of the intertemporal consumption model and provide a natural test of the plausibility of the fixed costs hypothesis as an explanation to limited participation by setting a lower bound on the level of costs needed to rationalize incomplete portfolios. The plausibility of the “cost of participation hypothesis” is an empirical issue and depends crucially on the size of the actual costs of financial market participation relative to the theoretical predictions: for the cost of participation hypothesis not to be rejected, the actual costs should be higher than my estimates. Overall, the empirical evidence on the nature and on the entity of the costs associated to financial transactions is rather limited, especially at the household level, and a reason is that some of these costs are likely to be figurative and related to information gathering and processing, especially when it comes to investing in the stock market. Given a consumption-income ratio of 0.5, spending just one hour per week in caring for one’s investment would add up to a non-monetary cost of 5 percent of household expenditure. To this one must add monetary charges, which based on the US National Income and Product Accounts could amount to up to 3 percent of household consumption. These figures suggest that the true costs of participating to the stock market are likely to be larger than my estimated bounds, supporting the claim that the heuristic test does not reject the “participation cost hypothesis”, even at a reasonably high degree of confidence, and this leads me to think that even more powerful tests would not be rejected. In fact, due to measurement error in consumption data which is a source of upward bias, the "true" bounds are likely to be smaller than the ones that I find. Hence, based on my heuristic test the hypothesis that participation costs may reconcile the stock market participation anomaly with the intertemporal choice model cannot be rejected. However, the heterogeneity in the gains from participation and the fact that the wealthiest non-participants appear to forego substantially larger returns suggest that some caution is due when considering the participation cost hypothesis to explain non-participation among those with substantially large amounts of financial wealth, which is in line with Vissing-Jørgensen (2003) conclusions. Finally, one cannot rule out that participation costs may reconcile also the choices of those households who hold neither risky, nor riskless assets. In fact, relatively modest costs are sufficient to rationalize their behaviour, given some constraint on the possibility of financing the investment in the risky asset by borrowing at the riskless rate to finance the investment in risky assets.

References

Attanasio, O., and G. Weber, 1995, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103, 1121-1157.

Attanasio, O., Banks, J., and S. Tanner, 2002, "Asset Holding and Consumption Volatility," *Journal of Political Economy*, 110, 771-792.

Cochrane, J., 1989, "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *American Economic Review*, 79, 319-337.

Constantinides, G., 1982, "Intertemporal Asset Pricing with Heterogeneous Consumers and without Demand Aggregation", *Journal of Business*, Vol. 55: 253-67.

Grossman, S., and R. Shiller, 1982, "Consumption Correlatedness and Risk Measurement in Economies with Non-Traded Assets and Heterogeneous Information," *Journal of Financial Economics*, 10, 195-210.

Hansen, L., and K. Singleton, 1982, "Generalized Instrumental Variable Estimation of Nonlinear Rational Expectation Models," *Econometrica*, 50, 1269-1286.

Heaton, J., and D. Lucas, 2000, "Portfolio Choice in the Presence of Background Risk," *Economic Journal*, 110, 1-26.

Heaton, J., and D. Lucas, 1997, "Market Frictions, Saving Behavior and Portfolio Choice," *Macroeconomic Dynamics*, 1, 76-101.

Kandel S. and R. Stambaugh, 1991, "Expectations and Volatility of Consumption and Asset Returns," *Review of Financial Studies*, 3, 207-232.

Kocherlakota, N., 1990, "On Tests of the Representative Agent Asset Pricing Models," *Journal of Monetary Economics*, 6, 285-304.

Lucas, D., 1994, "Asset Pricing with Undiversifiable Risk and Short Sales Constraints: Deepening the Equity Premium Puzzle," *Journal of Monetary Economics*, 34, 325-342.

Luttmer, E., 1999, "What Level of Fixed Costs Can Reconcile Asset Returns and Consumption Choices?," *Journal of Political Economy*, 107, 969-997.

Mankiw, N., and S. Zeldes, 1991, "The Consumption of Stockholders and Non-Stockholders," *Journal of Financial Economics*, 17, 211-219.

Mehra, R., and E. Prescott, 1985, "The Equity Premium Puzzle," *Journal of Monetary Economics*, 15, 145-161.

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

Newey, W., and K. West, 1987, "A Simple Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703-08.

Paiella, M., 2001, Limited Financial Market Participation: A Transaction Cost-based Explanation, Institute of Fiscal Studies Working Paper No. 0106.

Paiella, M., 2004, Revisiting the Implications of Heterogeneity in Financial Market Participation for the C-CAPM, *European Finance Review*, 8, 1-36.

Poterba, J., and A. Samwck, 1997, "Household Portfolio Allocation Over the Life Cycle," NBER Working Paper 6185.

Samuleson, P., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics*, 51, 39-46.

Simon, H., 1978, "Rationality as Process and as Product of Thought," *American Economic Review*, 68, 1-16.

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

Vissing-Jørgensen, A., 2003, "Perspectives on Behavioral Finance: Does 'Irrationality' Disappear with Wealth? Evidence from Expectations and Actions," NBER Macroeconomics Annual.

[1] An additional issue is related to the fact that, over the 80s and 90s, the proportion of households holding shares has increased dramatically. Hence, composition effects might be important in the sub-sample of shareholders.

[2] Similar figures obtain using the Survey of Consumer Finance.

[3] There is no real consensus as to the value that the coefficient of relative risk aversion should take on. Mehra and Prescott (1985) quote several microeconomic estimates that bound it from above by 3. Kandel and Stambaugh (1991) and Kocherlakota (1990) show that individuals' risk aversion might be higher. However, the current state of profession seems to share Lucas (1994)'s view that any proposed solution that does not explain the premium for a coefficient of relative risk aversion equal to 2.5 or less is likely to be viewed as a resolution that depends on a high degree of risk aversion.

[4] The Monte Carlo simulation is based on the assumption that measurement error is multiplicative and independent across households and from the level of consumption and asset returns and that is lognormally distributed with mean one. Its variance is computed as the difference between the variance of CEX consumption of non-participant and the variance of aggregate US expenditure, which is my proxy for measurement error-free consumption. The simulation suggests that, with log-utility, my estimator is biased upward by around 5 percent; the bias goes up to 20 percent when risk aversion is set equal to three, to 31 percent when it is equal to 8.

[5] The monetary charges-to-consumption estimate is based on NIPA Table 2.5.5, "Personal Consumption Expenditures by Type of Expenditure". The numerator has been computed by adding lines 61 (Brokerage charges and investment counselling), 62 (Bank service charges, trust service and safe deposit box rental, excluding commercial bank service charges and fees) and 63 (Services furnished without payment by financial intermediaries except life insurance carriers, excluding services by commercial banks). The denominator is total expenditure on non-durable goods and services.

[6] One hour per week corresponds roughly to one fortieth of income. In the sample used, c_t/y_t =0.5, where c_t denotes expenditure on non-durable goods and services and y_t is after-tax income. Then, y_t,= $c_t/0.5$ and one fortieth of income corresponds to (1/40) x (1/0.5) x c_t which is equal to 0.05 c_t.

[7] Households holding neither risky nor riskless assets are allowed to smooth consumption by means of other, non-financial assets and their consumption does not necessarily coincide with their endowment.

[8] These are first-order conditions, which are necessary, but not sufficient for a maximum, unless the function being maximized is strictly concave in the parameters, which needs not be the case in the problem considered here. Thus, second-order conditions must be checked as well.

[9] In principle, when bounding the costs of investing in risky assets I could consider also type 3 households, rather than only type 2. However, the foregone benefits for non-investing in risky assets for type 3 households are likely to be very different from those of type 2 households because ultimately the foregone gains depend also on whether the household has other assets. For type 2 households much of the gains of investing in risky assets can be expected to come from the exploitation of the equity premium. Instead for type 3 households, they can be expected to be related to consumption smoothing, although at a somewhat high risk. Averaging the two types of gains would make the estimated *d* little informative and more difficult to interpret.

[10] This timing implies a nine month gap between *t* and *t+1* and observations on consumption growth over nine months. Hence, for the estimation I use the return on investments of nine months. The bound to the gains from a

nine-month investment are multiplied by 4/3 to recover the bound to the gains from an investment of twelve months.

[11] Nominal consumption is deflated by means of household specific indices based on the Consumer Price Index provided by the Bureau of Labor Statistics. The individual deflators are determined as geometric averages of elementary regional price indices, weighted by the shares of household expenditure on individual goods. See Attanasio and Weber (1995) for a more extensive discussion of these indices. Household per-adult equivalent consumption is obtained from total household consumption using the following adult equivalence scale: the household head is weighted 1, the other adults and the children are weighted 0.8 and 0.4, respectively. Using per-adult equivalent consumption allows to account for changes in consumption deriving from changes in household composition occurred in-between interviews.

[12] Over the period considered stock returns were abnormally high. This implies that any foregone gain estimate based on data covering these years tend to overstate expected benefits.

[13] Higher rates imply slightly higher estimates of *d*, but the overall conclusions do not change in any significant respect.

[14] Results are available upon request.

[15] The gains of those with high wealth and low initial consumption are slightly lower if they are prevented from short-selling the risky asset, which allows them to smooth consumption over the two periods.

[16] $y_t(z_t; a_y) < -x^c(a_c)$ is never optimal as it would imply borrowing at the risky rate to invest in the riskless asset.

[17] See the Federal Reserve Board Statistical Release “G19, Consumer Credit”.

Table 1: Sample Composition

Year	Share of households with risky and riskless assets (%) (Type 1)	Share of households with just riskless assets (%) (Type 2)	Share of households with neither risky nor riskless assets (%) (Type 3)	Total households
1982	26.8	54.6	18.6	1,905
1983	27.3	54.1	18.5	1,943
1984	28.0	53.9	18.1	1,920
1985	25.6	53.8	20.6	944
1986	30.5	51.8	17.6	1,885
1987	31.5	50.8	17.7	1,902
1988	30.9	49.7	19.4	1,951
1989	30.6	50.2	19.3	1,945
1990	29.5	51.3	19.2	1,918
1991	34.7	45.0	20.3	1,976
1992	35.0	45.7	19.3	1,789
1993	33.2	47.0	19.8	1,856
1994	33.9	46.8	19.3	1,884
1995	38.2	40.8	21.1	152
Total	30.8	50.2	19.0	23,970

Note: The relatively small number of households in 1985 is due to the fact that in 1986 the sample was redesigned and many households who entered the survey in the second half of 1985 were dropped or had their identifier changed and I had to exclude them from my sample altogether because there is no information on their financial portfolios. As to 1995, only those households who enter the sample in January are included.

Table 2: Descriptive statistics for the total sample and for the three types of households

	Type 1	Type 2	Type 3	Whole Sample
Age*	43.9	42.3	43.7	43.1
Less than high school (%)	5.8	14.7	32.2	15.3
High school diploma (%)	51.1	59.3	53.4	56.7
College degree (%)	42.1	26.1	14.4	29.1
Male (%)	77.8	69.4	59.0	70.0
Single person (%)	14.3	20.1	18.0	17.9
Married (%)	76.8	63.9	53.6	65.9
Children (%)	49.2	47.8	53.5	49.3
Living in the Northeast (%)	21.3	18.3	28.1	21.1
Living in the Mid-West (%)	27.4	26.0	24.7	26.2
Living in the South (%)	26.8	28.2	29.8	28.1
Living in the West (%)	24.0	27.4	17.5	24.4
After tax annual household income*	$64,300	$46,400	$36,000	$49,900
Household financial wealth*	$52,100	$11,000	$0	$21,600
Annual household consumption *	$22,900	$18,200	$15,800	$19,200
No. of observations	7,388	12,021	4,561	23,970

Note: Income, financial wealth and consumption are in dollars of year 2000. Consumption consists of spending on non-durables and services. * denotes means.

Table 3: Annual Risky and Riskless Returns (1981:12-1995:09)

	Mean	Standard Deviation	Min	Max
S&P500 Composite Index	0.122	0.170	-0.277	0.759
3-month Treasury Bills	0.026	0.019	-0.007	0.076

Note: Real returns.

Table 4: The Foregone Gains for non-Investing in Risky Assets, with No Wealth Reallocation

Panel (a)

RRA	**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
x^c	0.072	0.050	0.042	0.037	0.034	0.033	0.030	0.027
	(0.009)	(0.004)	(0.003)	(0.003)	(0.004)	(0.006)	(0.011)	(0.015)
d	0.006	0.007	0.008	0.010	0.013	0.017	0.027	0.034
	(0.002)	(0.001)	(0.001)	(0.002)	(0.003)	(0.006)	(0.017)	(0.029)

Panel (b)

RRA	**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
average x^c_t	0.072	0.050	0.041	0.037	0.034	0.032	0.029	0.026
d	0.007	0.008	0.009	0.010	0.013	0.018	0.035	0.054
	(0.002)	(0.001)	(0.001)	(0.002)	(0.003)	(0.006)	(0.017)	(0.028)
$x_t^c(z_t;\alpha_c)$								
Constant	-0.601	-0.211	0.258	0.272	0.798	1.209	2.595	3.427
	(1.169)	(0.559)	(0.400)	(0.377)	(0.482)	(0.716)	(1.143)	(1.335)
R_t	-0.179	-0.080	-0.056	-0.050	-0.043	-0.043	-0.040	-0.005
	(0.057)	(0.028)	(0.023)	(0.024)	(0.028)	(0.036)	(0.057)	(0.102)
R^f_t	0.879	0.358	-0.121	-0.143	-0.642	-1.022	-2.312	-3.115
	(1.090)	(0.528)	(0.384)	(0.360)	(0.453)	(0.664)	(1.052)	(1.217)
term spread$_t$	2.178	1.214	0.934	1.260	1.118	1.252	1.161	0.934
	(1.194)	(0.486)	(0.377)	(0.460)	(0.601)	(0.748)	(0.905)	(1.216)
$(P/E)_t$	-0.365	-0.203	-0.274	-0.317	-0.472	-0.638	-1.119	-1.441
	(0.408)	(0.172)	(0.121)	(0.130)	(0.184)	(0.265)	(0.408)	(0.439)

Note: CRRA utility. RRA is the coefficient of relative risk aversion. d denotes the bound estimate as a fraction of time t consumption. The trading rule for consumption is specified as: $x_t^c(z_t;a_c) = a_c' z_t$. x^c and '*average* x^c_t' denote the (average) optimal share of time t consumption to invest in the risky asset when households do not and do time the market when determining their investment, respectively. The estimates in Panel (a) are based on the assumption that households do not time the market (z_t=1). The estimates in Panel (b) are based on the assumption that households predict the equity premium using past returns on the S&P500 CI (R_t) and on 3-month Treasury bills (R^f_t), the term spread and the price-earnings ratio. The discount rate is set equal to 0.98. The sample includes 12,021 households who do not own risky assets, but do own riskless ones. Standard errors in parentheses. To compute the standard errors, a Newey and West (1987) type of correction has been used to account for the MA(9) structure of the residual, due to the overlapping of the observations of the utility gains.

Table 5: The Foregone Gains for non-Investing in Risky Assets, with Wealth Reallocation

Panel (a)

RRA	**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
x^w	1.00	0.85	0.65	0.50	0.40	0.35	0.25	0.10
x^c	0.049	0.037	0.033	0.031	0.030	0.030	0.029	0.027
	(0.005)	(0.002)	(0.002)	(0.003)	(0.004)	(0.007)	(0.011)	(0.015)
d	0.033	0.022	0.018	0.017	0.017	0.020	0.028	0.034
	(0.010)	(0.008)	(0.005)	(0.004)	(0.004)	(0.006)	(0.017)	(0.029)

Panel (b)

RRA	**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
x^w	1.00	0.85	0.65	0.50	0.40	0.35	0.25	0.25
average x^c_t	0.050	0.037	0.033	0.031	0.030	0.029	0.027	0.025
d	0.033	0.022	0.018	0.017	0.018	0.021	0.036	0.054
	(0.010)	(0.008)	(0.005)	(0.004)	(0.004)	(0.006)	(0.017)	(0.028)
$x_t^c(z_t;\alpha_c)$								
Constant	-0.387	0.385	0.322	0.464	0.823	1.456	2.786	3.505
	(0.638)	(0.257)	(0.278)	(0.361)	(0.535)	(0.797)	(1.268)	(1.533)
R_t	-0.083	-0.017	-0.019	-0.026	-0.034	-0.037	-0.046	-0.012
	(0.033)	(0.016)	(0.019)	(0.024)	(0.031)	(0.039)	(0.061)	(0.110)
R^f_t	0.535	-0.291	-0.236	-0.360	-0.683	-1.267	-2.487	-3.181
	(0.592)	(0.237)	(0.260)	(0.340)	(0.501)	(0.741)	(1.171)	(1.397)
term spread$_t$	1.109	0.580	0.696	0.943	1.161	1.063	1.200	1.012
	(0.629)	(0.378)	(0.426)	(0.520)	(0.674)	(0.832)	(1.020)	(1.378)
$(P/E)_t$	-0.211	-0.244	-0.224	-0.303	-0.458	-0.669	-1.165	-1.464
	(0.217)	(0.102)	(0.116)	(0.148)	(0.208)	(0.296)	(0.436)	(0.480)

Note: See note to table 4. x^w denotes the optimal share of wealth invested in the riskless asset to be moved into the risky asset and is determined by grid search.

Table 6: The Foregone Gains for Different Groups of Non-Stockholders, with No Wealth Reallocation

RRA		**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
Low wealth and low consumpt. (1)	*average* x^c_t	0.062	0.045	0.041	0.039	0.039	0.038	0.030	0.017
	d	0.005	0.006	0.006	0.006	0.003	0.000	0.009	0.039
		(0.001)	(0.001)	(0.001)	(0.003)	(0.006)	(0.012)	(0.032)	(0.046)
Low wealth and high consumpt. (2)	*average* x^c_t	0.079	0.059	0.052	0.048	0.044	0.037	0.024	0.013
	d	0.009	0.010	0.010	0.006	0.000	0.000	0.001	0.007
		(0.002)	(0.003)	(0.004)	(0.007)	(0.010)	(0.013)	(0.017)	(0.017)
Intermed. wealth and low consumpt. (3)	*average* x^c_t	0.064	0.045	0.039	0.035	0.033	0.030	0.023	0.017
	d	0.006	0.006	0.005	0.003	0.002	0.004	0.019	0.033
		(0.002)	(0.002)	(0.002)	(0.003)	(0.005)	(0.007)	(0.015)	(0.028)
Intermed. wealth and high consumpt. (4)	*average* x^c_t	0.080	0.059	0.052	0.047	0.045	0.041	0.029	0.022
	d	0.009	0.011	0.012	0.011	0.004	0.000	0.009	0.029
		(0.002)	(0.002)	(0.003)	(0.005)	(0.009)	(0.013)	(0.020)	(0.023)
High wealth and low consumpt. (5)	*average* x^c_t	0.057	0.037	0.031	0.028	0.024	0.019	0.008	0.001
	d	0.004	0.004	0.002	0.000	0.000	0.010	0.047	0.076
		(0.001)	(0.001)	(0.002)	(0.003)	(0.005)	(0.008)	(0.014)	(0.017)
High wealth and low consumpt. (6)	*average* x^c_t	0.080	0.059	0.050	0.044	0.039	0.034	0.024	0.014
	d	0.009	0.011	0.012	0.012	0.009	0.006	0.013	0.049
		(0.002)	(0.003)	(0.004)	(0.008)	(0.019)	(0.041)	(0.141)	(0.258)

Note: See note to table 4. (1) The sample consists of 2,726 households whose financial wealth is in the lowest third of the sample financial wealth distribution and whose consumption is below the sample median. (2) The sample consists of 1,293 households whose financial wealth is in the lowest third of the distribution and whose consumption is above the median. (3) The sample consists of 1,957 households whose financial wealth is in the middle third of the distribution and whose consumption is below the median. (4) The sample consists of 2,005 households whose financial wealth is in the middle third of the distribution and whose consumption is above the median. (5) The sample consists of 1,379 households whose financial wealth is in the top third of the distribution and whose consumption is below the median. (6) The sample consists of 2,661 households whose financial wealth is in the top third of the distribution and whose consumption is above the median.

Table 7: The Foregone Gains for Different Groups of Non-Stockholders, with Wealth Reallocation

RRA		**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
Low wealth and low consumpt. (1)	x^w	1.00	1.00	1.00	1.00	1.00	0.00	0.00	0.00
	average x^c_t	0.060	0.043	0.039	0.037	0.038	0.038	0.030	0.017
	d	0.007	0.008	0.008	0.008	0.004	0.000	0.009	0.039
		(0.002)	(0.001)	(0.001)	(0.002)	(0.005)	(0.012)	(0.032)	(0.046)
Low wealth and high consumpt. (2)	x^w	1.00	1.00	1.00	1.00	0.00	0.00	0.00	0.00
	average x^c_t	0.077	0.058	0.050	0.046	0.044	0.037	0.024	0.013
	d	0.011	0.012	0.011	0.007	0.000	0.000	0.001	0.007
		(0.003)	(0.003)	(0.004)	(0.007)	(0.010)	(0.013)	(0.017)	(0.017)
Intermed. wealth and low consumpt. (3)	x^w	1.00	1.00	1.00	1.00	1.00	1.00	0.75	0.45
	average x^c_t	0.051	0.032	0.027	0.025	0.025	0.024	0.021	0.018
	d	0.019	0.018	0.017	0.013	0.010	0.010	0.021	0.033
		(0.005)	(0.005)	(0.005)	(0.006)	(0.006)	(0.007)	(0.012)	(0.026)
Intermed. wealth and high consumpt. (4)	x^w	1.00	1.00	0.95	0.65	0.35	0.25	0.00	0.00
	average x^c_t	0.073	0.052	0.045	0.042	0.042	0.039	0.029	0.022
	d	0.016	0.017	0.018	0.016	0.006	0.001	0.009	0.029
		(0.004)	(0.004)	(0.004)	(0.006)	(0.009)	(0.013)	(0.020)	(0.023)
High wealth and low consumpt. (5)	x^w	0.65	0.55	0.55	0.40	0.35	0.15	0.05	0.05
	average x^c_t	-0.001	-0.003	-0.001	0.000	-0.002	0.001	0.000	-0.007
	d	0.067	0.046	0.036	0.030	0.032	0.037	0.060	0.087
		(0.020)	(0.016)	(0.015)	(0.011)	(0.011)	(0.009)	(0.014)	(0.017)
High wealth and low consumpt. (6)	x^w	1.00	0.85	0.55	0.40	0.30	0.25	0.25	0.00
	average x^c_t	0.031	0.030	0.029	0.027	0.025	0.022	0.013	0.014
	d	0.065	0.042	0.031	0.025	0.020	0.013	0.021	0.049
		(0.021)	(0.018)	(0.013)	(0.013)	(0.020)	(0.036)	(0.104)	(0.258)

Note: see note to table 6. x^w denotes the optimal share of wealth invested in the riskless asset to be moved into the risky asset and is determined by grid search.

Table 8: The Foregone Gains for non-Investing in Riskless Assets and the Corresponding Optimal Portfolios

Panel (a)

RRA		**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
	x^c	0.052	0.048	0.045	0.042	0.040	0.036	0.029	0.024
	d_{rf}	0.003	0.007	0.011	0.016	0.025	0.035	0.048	0.048
		(0.001)	(0.001)	(0.002)	(0.005)	(0.011)	(0.022)	(0.045)	(0.052)
$x^c(\alpha_c)$									
	Constant	2.907	2.993	3.062	3.119	3.184	3.273	3.495	3.702
		(0.235)	(0.128)	(0.112)	(0.145)	(0.231)	(0.351)	(0.570)	(0.685)

Panel (b)

RRA		**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
	average x^c_t	0.052	0.048	0.045	0.042	0.038	0.033	0.028	0.032
	d_{rf}	0.003	0.007	0.011	0.016	0.026	0.038	0.058	0.063
		(0.001)	(0.001)	(0.002)	(0.005)	(0.012)	(0.024)	(0.048)	(0.056)
$x_t^c(z_t;\alpha_c)$									
	Constant	1.509	1.639	1.880	1.857	-11.677	-26.752	-116.934	-169.754
		(15.091)	(7.607)	(6.772)	(9.070)	(13.218)	(33.190)	(96.316)	(112.778)
	R_t	0.090	0.109	-0.111	-0.386	-0.684	-0.421	2.373	3.663
		(1.675)	(0.884)	(0.758)	(0.855)	(1.399)	(2.356)	(3.165)	(3.528)
	R^f_t	1.274	1.207	1.272	1.645	15.312	29.998	116.528	167.595
		(15.155)	(7.342)	(6.518)	(8.891)	(12.986)	(31.721)	(93.699)	(109.852)

Panel (c)

RRA		**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
	average x^c_t	0.052	0.048	0.045	0.041	0.037	0.031	0.015	0.004
	d_{rf}	0.003	0.007	0.011	0.017	0.026	0.038	0.058	0.066
		(0.001)	(0.001)	(0.002)	(0.005)	(0.012)	(0.024)	(0.046)	(0.050)
$x_t^c(z_t;\alpha_c)$									
	Constant	0.140	0.229	0.314	0.471	0.675	0.896	1.266	1.476
		(0.813)	(0.389)	(0.306)	(0.352)	(0.458)	(0.555)	(0.564)	(0.422)
	R_t	-0.002	0.002	0.018	0.039	0.065	0.096	0.177	0.252
		(0.088)	(0.045)	(0.039)	(0.049)	(0.069)	(0.098)	(0.148)	(0.136)
	R^f_t	-0.084	-0.180	-0.282	-0.461	-0.692	-0.947	-1.411	-1.706
		(0.832)	(0.389)	(0.307)	(0.364)	(0.485)	(0.601)	(0.647)	(0.421)

Note: CRRA utility. RRA is the coefficient of relative risk aversion. d_{rf} denotes the bound estimate as a fraction of time t consumption. x^c and '*average* x^c_t' denote the (average) optimal share of time t consumption to invest in the riskless asset when households do not and do time the market when determining their investment, respectively. The estimates in Panel (a) are based on the assumption that households do not time the market when determining their investment and cannot short-sell the riskless asset. The (constant) trading rule for consumption is specified as: $x^c(a_c) = (1 + \exp(a_c))^{-1}$. The estimates in Panel (b) are based on the assumption that households do time the market using the past returns on the S&P500 CI (R_t) and on 3-month Treasury bills (R^f_t). The trading rule for consumption is specified as $x_t^c(z_t; a_c) = (1 + \exp(a_c' z_t))^{-1}$. The estimates in Panel (c) are based on the assumption that households time the market and are allowed to short-sell the riskless asset. The trading rule for consumption is specified as: $x_t^c(z_t; a_c) = a_c' z_t$. The discount rate is set equal to 0.98. The sample includes 4,561 households who own neither risky nor riskless assets.

Table 9: The Foregone Gains for non-Investing in Either Risky or Riskless Assets and the Corresponding Optimal Portfolios

Panel (a)

RRA		**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
	$(x^c + y)$	0.075	0.056	0.049	0.043	0.039	0.035	0.027	0.000
	$(-y)$	0.000	0.001	0.001	0.001	0.001	0.001	0.002	0.024
	$d_{r,rf}$	0.007	0.010	0.013	0.019	0.027	0.036	0.049	0.048
		(0.001)	(0.001)	(0.002)	(0.005)	(0.011)	(0.022)	(0.045)	(0.052)
$x^c(\alpha)$									
	Constant	2.516	2.811	2.960	3.068	3.159	3.267	3.493	3.703
		(0.086)	(0.066)	(0.076)	(0.120)	(0.211)	(0.339)	(0.570)	(0.688)
$y(\alpha_y)$									
	Constant	8.341	7.458	7.261	6.707	6.628	6.576	6.102	3.703
		(0.026)	(0.025)	(0.025)	(0.023)	(0.020)	(0.017)	(0.015)	(0.017)

Panel (b)

RRA		**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
	$(x^c + y)$	3.019	1.303	0.902	0.633	0.451	0.312	0.097	0.000
	$(-y)$	-3.010	-1.279	-0.877	-0.603	-0.417	-0.278	-0.068	0.024
	$d_{r,rf}$	0.127	0.057	0.042	0.036	0.037	0.041	0.049	0.048
		(0.061)	(0.026)	(0.018)	(0.013)	(0.013)	(0.022)	(0.045)	(0.052)
$x^c(\alpha)$									
	Constant	4.754	3.727	3.627	3.468	3.368	3.352	3.496	3.703
		(0.017)	(0.593)	(0.419)	(0.307)	(0.322)	(0.403)	(0.580)	(0.688)
$y(\alpha_y)$									
	Constant	3.010	1.279	0.877	0.603	0.417	0.278	0.068	3.703
		(0.001)	(0.198)	(0.255)	(0.211)	(0.170)	(0.147)	(0.164)	(0.017)

Panel (c)

RRA		**1**	**2**	**3**	**4**	**5**	**6**	**8**	**10**
	$(x^c + y_1)$	1.153	0.534	0.318	0.193	0.092			
	$(-y_1)$	-1.087	-0.483	-0.271	-0.149	-0.052		No borrowing	
	$d_{r,rf}$	0.018	0.014	0.016	0.020	0.027			
		(0.022)	(0.010)	(0.006)	(0.005)	(0.011)			
$x^c(\alpha)$									
	Constant	2.648	2.910	3.022	3.091	3.163			
		(0.111)	(0.102)	(0.107)	(0.142)	(0.222)			
$y_1(\alpha_{y1})$									
	Constant	1.087	0.483	0.271	0.149	0.052			
		(0.963)	(0.512)	(0.317)	(0.216)	(0.161)			

Note: CRRA utility. RRA denotes the coefficient of relative risk aversion. $d_{r,rf}$ denotes the bound estimate as a fraction of time *t* consumption. $(x^c + y)$ and $(x^c + y_1)$ denote the optimal share of time *t* consumption to invest in the risky asset. $(-y)$ denotes the optimal share of time *t* consumption to invest in the riskless asset. When $(-y) \leq 0$ borrowing occurs. $(-y_1)) \leq 0$ indicates borrowing at a rate equal to twice the riskless rate. The trading rule for consumption is specified as: $x^c(a_c) = (1 + \exp(a_c))^{-1}$. The estimates in Panel (a) are based on the assumption that households cannot borrow at the riskless rate to finance the risky asset investment and the investment in the riskless asset is specified as $(-y(a_y)) = (1 + \exp(a_y))^{-1}$. The estimates in Panel (b) are based on the assumption that they can borrow as much as they want at the riskless rate and $y(a_y) = a_y$. The estimates in Panel (c) are based on the assumption that households can borrow at a rate equal to twice the riskless rate and $y_1(a_{y_1}) = a_{y_1}$.