ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Paiella, Monica

Working Paper

Limited financial market participation: A transaction cost-based explantation

IFS Working Papers, No. 01/06

Provided in Cooperation with:
Institute for Fiscal Studies (IFS), London

Suggested Citation: Paiella, Monica (2001) : Limited financial market participation: A transaction cost-based explantation, IFS Working Papers, No. 01/06, Institute for Fiscal Studies (IFS), London, https://doi.org/10.1920/wp.ifs.2001.0106

This Version is available at:
https://hdl.handle.net/10419/71530

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

LIMITED FINANCIAL MARKET PARTICIPATION: A TRANSACTION COST-BASED EXPLANATION

Monica Paiella

THE INSTITUTE FOR FISCAL STUDIES
WP 01/06

Limited Financial Market Participation: A Transaction Cost-based Explanation

Monica Paiella[1]

University College London,
Bank of Italy and
Institute for Fiscal Studies

15 February 2001

Abstract

This paper focuses on the issue of limited financial market participation and determines a lower bound on the level of fixed transaction costs that are required to reconcile observed portfolio choices with asset returns within an isoelastic utility framework. The bound is determined from the set of conditions that ensure the optimality of consumption behavior by financial market non-participants. It represents the lowest possible cost rationalizing observed non-participation choices by providing a measure of the forgone utility gains from participation for observed non-participants. Such gains are related both to the magnitude of financial market returns and to the opportunity of smoothing consumption, with the benefits of the former decreasing in the degree of relative risk aversion and those of the latter increasing in it. Using the US Consumer Expenditure Survey, I find that a yearly cost of at least $70 is needed to rationalize non-participation for a consumer with log utility and who can trade in the S&P500 CI. This lower bound declines rapidly in risk aversion for levels of risk aversion up to two/three; for higher values, it levels off. A yearly cost of at least $31 is needed to rationalize non-participation for a consumer with log utility and who can trade in US Treasury Bills. This lower bound rises steadily in risk aversion.

JEL Classification: G11, D12 and E21
Keywords: household portfolio allocation, financial market participation, transaction costs, heterogeneous consumers

---

[1] E-mail: monicapaiella@tin.it. Thanks are due to Professor Orazio Attanasio for numerous insights and suggestions and to Professor Richard Blundell and Professor Costas Meghir for helpful discussion. Also, I thank for their comments James Banks, Marco Cagetti, Janice Eberly, Hamish Low, Erzo Luttmer, Michael Haliassos, Luigi Pistaferri, Andrea Tiseno, Annette Vissing-Jørgensen, Joachim Winter, IFS internal seminars participants and participants to the 1999 NBER Summer Workshop on the Aggregate Implications of Microeconomic Behavior and to the 1999 TMR conference on Savings, Pensions and Portfolio Choice. Financial support from the Institute for Fiscal Studies and the EEC (TMR, grant n. ERBFMBICT 972537) is gratefully acknowledged. The views expressed are those of the author and do not necessarily reflect those of the Bank of Italy. Any errors are responsibility of the author.

1. Introduction

A large number of studies has suggested that observed asset returns are inconsistent with consumption choices as predicted by the standard neo-classical model for consumption. The testable implications of this model have, in fact, repeatedly been proven to be at odds with empirical evidence and have given rise to the equity premium and other asset pricing puzzles[2]. Such empirical inconsistency has generally been rationalized by the literature either assuming that agents are highly averse to consumption risk or conjecturing that trading stock is much more costly than trading bonds. Recently, it has also been shown that accounting for limited stock market participation might be important for explaining the puzzles, since allowing for differences in the consumption patterns of asset holders and non-holders tends to lower the risk aversion implied by the model[3]. However, no attempt has been made to rationalize non-participation. Non-participation to financial markets is the main issue this paper wants to address and does so by verifying whether it can be rationalized on the ground of transaction costs that are small enough to be *realistic*. The second issue the paper deals with is that of the differences in the costs of trading distinct assets. In the literature, cost differentials generally result from calibration exercises, whereas here I identify the bounds to the costs directly and look for evidence that trading risky assets is costlier than trading riskless ones.

The approach adopted to identify the transaction costs is based on the observation that the standard way of examining the consistency of a model with the empirical evidence is to test a set of first-order conditions against the data. The rejection of such conditions suggests that there are gains the consumer could make by modifying her consumption. However, if such gains are not too large, a possible interpretation of the sub-optimal behavior is that the consumer faces small transaction costs every time she approaches financial markets and the costs of modifying consumption are higher than the utility gains. By measuring such gains it is possible to determine a set of bounds on the level of transaction costs that help rationalize non-participation and, ultimately, can reconcile asset returns and consumption choices.

[2] See Kocherlakota (1996) for a thorough review of the literature.
[3] See Attanasio, Banks and Tanner (1996) for a study based on the UK Family Expenditure Survey and Paiella (1999) and Vissing-Jørgensen (1999) for two analyses based on the US Consumer Expenditure Survey.

For the estimation of such cost bounds, I extend Luttmer (1999) and determine the lower bounds as the minimal costs that rationalize non-participation, i.e. as those costs exactly equal to the utility gains from trade. However, unlike Luttmer, whose work is based on aggregate information, I use individual level data, which allow to distinguish between actual participants and non-participants to financial markets, instead of simply characterizing traders and non-traders in the time period under scrutiny. As a consequence, the nature of the costs I focus upon is substantially different from the nature of the costs in Luttmer's analysis. In fact, the frictions he considers are the costs that the representative agent must pay to trade and modify her consumption in the current period and in one or at most few subsequent periods. Instead, by distinguishing between participants and non-participants, this paper focuses on the costs any individual faces in order to actually participate to financial markets. In addition, because of the use of aggregate data, the validity of Luttmer's results is limited and his analysis applies strictly only to an agent who happens to consume US per-capita consumption because, in the presence of fixed costs, the conditions upon which aggregation results are based do not hold. For this reason, the use of micro data is particularly desirable in a framework where fixed costs play a role. The use of individual-level data brings about several other advantages. First, it allows to verify whether there are important cost differences when trading different portfolios, - at least to the extent that the data permit to distinguish between different assets. Second, it allows to take into account the effects that individual specific factors have on utility reducing the scope for unobserved heterogeneity and, consequently, the potential for bias. Last, given the availability of some panel dimension in the data I use, it is possible to account for differences in the covariance between individual consumption growth and asset returns.

Another empirical paper studying the interaction between market frictions and household portfolio choice with micro data is Vissing-Jørgensen (1999). Vissing-Jørgensen (1999) is built on the methodology of Mulligan and Sala-i-Martin (1996) and differs substantially from my type of analysis. In fact, the objective of my work is to determine the minimal costs rationalizing the choice of holding no equity despite the premium and I find that relatively small costs can indeed justify such behavior. Instead, Vissing-Jørgensen (1999) uses a dynamic sample selection model to gather evidence of state dependency in financial market participation - which is symptomatic of entry and transaction costs - and a censored regression model to determine the

distribution of the per-period participation costs. She estimates the median of this cost to be around \$200[4], which is a higher figure than the ones I obtain, but is fully consistent with my results.

The costs I consider in the paper are fixed per-period participation costs that must be paid at the time of investment and in each subsequent period as long as the agent stays in the market. Since I estimate the bounds as foregone utility gains of non-participation, the costs I set limits upon can include both cash outlays and "figurative" charges, such as brokerage fees and other commissions, bid-ask spreads, money/time spent understanding financial markets and determining the optimal portfolio, money/time spent setting up and managing the accounts, value of time spent trading and any other kind of opportunity cost of investors' time in processing information.

The rest of the paper is organized as follows. In Section 2, I discuss the model for the gains from financial market participation and relate such gains to the trading costs. In Section 3, I examine the econometric issues arising from the estimation of the model and present the estimation procedure. In Section 4, I describe the data and analyze the empirical results. Section 5 concludes.

2. Measuring Transaction Costs

Consider an environment where households have access to several means to substitute consumption over time. In particular, they can accumulate real assets, currency and/or financial securities. The securities can be traded after the payment of a fixed cost that can vary between the market for risky assets and the market for riskless ones. Households have additively separable preferences over consumption and the per-period utility function is strictly increasing and concave. Let $\{c^h\}_t, t = 1,2,...$ be household h observed sequence of consumption choices. These choices are the solution to some unobservable maximization problem involving labour supply, saving and portfolio composition. On the basis of portfolio composition, it is possible to distinguish among three types of households: those who hold both risky and riskless assets (type 1); those who hold only riskless assets (type 2); and those who have chosen not to participate to any financial market (type 3). Households are utility

[4] Vissing-Jørgensen's figures are in 1982-84 dollars. The purchasing power of 200 dollars of 1982-84 corresponds approximately to 350 dollars of year 2000.

maximizers. As a consequence, since at time t they could have chosen any other sequence of consumption bundles, their time t expected gain from deviating from $\{c^h\}_t$ must be negative. In particular, for those households who have chosen not to participate to some or both financial markets, time t expected utility gain, $y^*_{h,t}$, from adopting an alternative saving/consumption strategy involving participation, must be non-positive, i.e.

$$y^*_{h,t} = E[v^i_{h,t+1}(x,\delta,u) \mid I_{h,t}] \leq 0, \qquad i = 2,3; \qquad \forall t \tag{1}$$

where $v^i_{h,t+1}(.)$ is the utility gain that type i household h can obtain by deviating from the observed sequence of consumption choices, $\{c^h\}_t$. Under the assumption of additively separable preferences, the utility gain of type i household h can be written in the following way:

$$v^i_{h,t+1}(.) = \left\{U(\tilde{c}^i_{h,t}(x,\delta)) - U(c_{h,t})\right\}\exp(u_{h,t}) + \beta\left\{U(\tilde{c}^i_{h,t+1}(x,\delta)) - U(c_{h,t+1})\right\}\exp(u_{h,t+1}) \tag{2}$$

$U(c_{h,t})$ is the utility from the level of consumption that has been chosen. $U(\tilde{c}^i_{h,t}(x,\delta))$ is the utility in case of optimal participation to the financial market(s) that type i household h has chosen to stay out of. Participation implies paying the fixed cost δ and holding the optimally determined portfolio x of securities. $\{\tilde{c}^i_h(x,\delta)\}_t$ denotes consumption in case of participation. β is a positive subjective discount rate and $u_{h,t}$ is an unobservable random taste shifter which captures individual heterogeneity. $u_{h,t}$ represents all the unobservable and unaccounted for factors that affect individual portfolio choices and that I do not explicitly model or control for. Specifically, within the framework defined by (1) and (2), it captures all those unobservable features of individual preferences that influence the financial market participation choice and therefore determine the size of the loss from deviating from $\{c^h\}_t$. $E[.\mid I_{h,t}]$ is household h expectation conditional on the information available at time t. (1) and (2) imply that, at time t, given the information available, financial market non-participants should not be able to pay the fixed cost, participate optimally to the market(s) they have chosen to stay out of and obtain a higher level of utility. Inequalities like (1) must hold for any t and $t+s$, $s \geq 1$[5].

[5] Focusing the analysis on two adjacent periods is not restrictive, as long as the per-period costs of participating to the market for one-period securities and to the market for n-period securities are the

The inequality in (1) does not allow to identify the fixed cost parameters that would reconcile observed consumption choices with the assumption that agents are rational and choose optimally. However, if the instantaneous utility function is strictly concave, $v^i_{h,t+1}(x,\delta,u)$ is strictly decreasing in the fixed costs δ. Then, I can replace δ with $d \leq \delta$ and the inequalities in (1) with equalities and look for lower bounds to the costs. Such lower bound coincides with the level of participation costs that would make the utility in case of participation exactly equal to the utility in case of non-participation. In other words, it coincides with the levels of costs offsetting exactly the gain from participation.

Two issues are worth discussing at this stage. The first relates to the benchmark I use to quantify the gain/loss from participation. As I have mentioned in the Introduction, the model is motivated by the desire of rationalizing observed behavior as optimal, despite the empirical inconsistency of the neo-classical model for consumption noted by Hansen and Singleton (1983), Mehra and Prescott (1985) and Hansen and Jagannathan (1991)[6], among others. The fixed cost bounds are essentially measures of the benefit from participation and their interpretation is straightforward: the lower the bounds, the smaller the expected utility gain from participation and, consequently, the lower the transaction costs needed to make participation disadvantageous and, therefore, non-participation *rational*. Thus, for this exercise to be interesting, I must determine the individual optimal investment in those assets that households do not hold and compare the utility associated to such portfolio with the utility associated to the choice made, *ceteris paribus*. Then, for those who hold only riskless assets, let $\tilde{c}^2_{h,t} = c_{h,t} - \delta_2 - x_{2,h,t}q_2$ and $\tilde{c}^2_{t+1} = c_{h,t+1} + x_{2,h,t}R_{2,t,t+1}$ denote time t and time $t+1$ consumption in case of participation to the market for such asset. For those who have chosen not to hold any financial assets, let $\tilde{c}^3_{h,t} = c_{h,t} - \delta_{12} - x_{1,h,t}q_1 - x_{2,h,t}q_2$ and $\tilde{c}^3_{h,t+1} = c_{h,t+1} + x_{1,h,t}R_{1,t,t+1} + x_{2,h,t}R_{2,t,t+1}$ denote time t and time $t+1$ consumption in case of participation to the markets for risky and

same. In this instance, by an arbitrage argument, the one period returns on the two assets must be the same.

[6] The three studies mentioned above characterize the inconsistency of the theory with the data in different ways. Hansen and Singleton (1983) reject the overidentifying restrictions of the model. Mehra and Prescott (1985) point out an *equity premium puzzle*. Hansen and Jagannathan (1991)

riskless assets. δ_2 is the fixed cost for the market for risky assets. δ_{12} is the joint cost of participating to both financial markets[7]. Given the participation cost, x_1 and x_2 are the individual optimal holdings of riskless and risky assets with time t prices q_1 and q_2 and time $t+1$ payoffs $R_{1,t,t+1}$ and $R_{2,t,t+1}$, respectively. As it will be shown in the next section, the optimal portfolios are determined by exploiting the fact that asset returns are to some extent predictable using a pricing kernel based on investors' utility. The specification adopted for $\tilde{c}^i_{h,t}$ and $\tilde{c}^i_{h,t+1}$ is a simplification and implies that the resources to be invested in the market subject to a cost are obtained by reallocating expenditure over time without modifying saving, whatever form it takes. Yet, since financial assets involving higher costs carry on average also higher returns, it is reasonable to expect that after paying the cost and investing in the higher return asset, the investor moves into this asset some of her wealth accumulated in other lower return assets. Alternatively, the income effect from higher returns might induce her to increase her consumption at time t, reducing overall savings[8]. The specification I have adopted has the advantage of not requiring the computation of household cash-in-hand, which is not directly available from the data I use. However, the inability to use individual cash-in-hand can be expected to bias downward the estimates, especially in the case where some asset is not subject to transaction costs. However, as it will be shown, information available on household after-tax income allows to quantify the importance of this bias.

The second issue worth mentioning relates to the nature of the costs of financial market participation *vis á vis* the fact that the analysis focuses explicitly only on two time periods and neglects any continuation payoff. The focus on only two

---

determine a set of bounds on the first two moments of a generic stochastic asset-pricing factor and find that the moments of the marginal rate of substitution are inconsistent with such bounds.

[7] Given the nature of the costs, (1) and (2) for type 3 households do not allow to identify the participation cost to each individual market separately, but only a single cost that pertains to both markets jointly.

[8] Consider for simplicity a household that holds its savings in a zero return costless asset. The budget constraints for time t and $t+1$ can be written as: $c_t = y_t + s_{t-1} - s_t$ and $c_{t+1} = y_{t+1} + s_t - s_{t+1}$, where c_t, y_t and s_t are time t consumption, income and saving, respectively. If households were allowed to reshuffle their savings when participating hypothetically to financial markets, then $\tilde{c}_t$ and $\tilde{c}_{t+1}$ could be defined as: $\tilde{c}_t = c_t + s_t - x_t - \delta$ and $\tilde{c}_{t+1} = c_{t+1} - s_t + x_t R_{t,t+1}$, where $\tilde{c}_t$ is consumption in case of financial market participation; $c_t + s_t$ is time t cash-in-hand, i.e. it is the amount of resources available for either consumption or investment at time t; x_t is the optimal portfolio of costly assets with return $R_{t,t+1}$ and δ is the per-period participation cost. The "simpler" specification I have adopted is

time periods can be justified by assuming that households are at an optimum conditional on the presence of transaction costs. The counterfactual implies switching consumption between the two periods under scrutiny, leaving everything else unchanged (at the optimum) and consequently I do not need to keep into account any other date. For this to hold, costs must be fixed and per-period. In principle, financial market participation involves three types of costs: an entry cost, a transaction/trading cost and a per-period participation cost. Entry costs consist in the time and money spent determining the household optimal portfolio and, to most extent, are likely to be fixed. Trading costs are likely to have a fixed component, consisting in commissions and in the value of time spent trading, and a variable one, proportional to the amount traded, related to bid-ask spreads and to commissions variable components. Finally, the per-period participation costs represent all the portfolio management monetary and opportunity costs. The different types of costs are likely to affect participation choices in different ways. In particular, when entry costs are present, the number of periods that households expect to stay in the market becomes crucial in determining investment choices. Similarly, when trading costs exist, the length of the investment is a crucial factor. Finally, in the presence of per-period costs, the length of the investment and/or of participation is irrelevant only if asset returns are assumed to be exogenous and, therefore, independent on the number of financial market participants, which in turn depends on the costs. Reasonably, all three types of costs can be expected to exist. The assumption of fixed per-period participation costs together with the focus on one period participation and investment can cause the actual costs to be somewhat underestimated if the entry and the trading costs are the most significant cost component and/or household investments are very long term. In fact, in this instance the actual gains from participation would be larger than the one estimated with the model in (1) and (2). The empirical evidence on the nature of the costs and on households movements in and out of financial markets is rather scarce. Yet, as to the first issue, the wide availability of low cost mutual funds is believed to have reduced effectively the costs of buying and trading a well-diversified portfolio. As to the second, using portfolio choice data from the 1984, 1989 and 1994 waves of the Panel Study of Income Dynamics, Vissing-Jørgensen (1999) finds widespread

---

primarily dictated by data limitations. In fact, it does not require the computation of s_t, which is not directly available from the data.

movements in and out of financial markets, with many households participating in one year but not the others. Such observed behavior suggests that household investments are rather short term and is consistent with the view that entry and trading costs are limited. Thus, the main cost components of financial market participation are likely to be portfolio management costs, related both to the time and money spent determining the optimal asset portfolio and to the time and money spent following financial markets, in order to form expectations on future returns and change the investment accordingly. If this is the case, although the former of these costs are likely to be somewhat higher for first-time investor (but not necessarily for new entrants), the assumption of fixed per-period participation costs should not cause the bound underestimation to be serious.

3. Estimation Issues

3.1. Econometric Issues

As explained in Section 2, after replacing the inequality in (1) with an equality, the model

$$E_t[v^i_{h,t+1}(x,d,u)] = 0 \qquad i = 2,3; \quad \forall t \tag{3}$$

allows to identify and measure a bound to the cost of financial markets participation. Such measure is provided by the value of d equalizing the expected utility from planned consumption to the expected utility from consumption in case of participation to some additional financial market, whose participation costs want to be quantified. One way of interpreting the fixed cost bound d is in terms of Hicks compensating variation for a change in prices from the set of (unobservable) prices implicit in the individual preferences to observable market prices. As a consequence, the cost bounds are in principle heterogeneous. Because of the lack of a long panel dimension in the data used, it is not possible to estimate consistently the bounds at the household level. However, I can compute an average individual household expected gain that will yield an estimate of the lower bound to the transaction costs for a consumer with a mean expected gain. Such estimate will differ from the mean of the individual lower bounds for a Jensen inequality term due to the fact that the utility gain function is strictly decreasing and concave in the cost. The issue can be illustrated in the following way. Assume that there are just two kinds of households. For the first, the expected gain from financial market participation is set to zero by a cost equal to d_1.

For the second, the expected gain is set to zero by a cost d_2 , with $d_2 > d_1$. The mean of the expected gains (and consequently the mean of the costs), $\bar{d}$, is simply the average of d_1 and d_2. Due to the inability of identifying the individual expected gains, I cannot determine $\bar{d}$, but I can look for the bound to the cost for a consumer whose expected utility gain coincides with the households mean expected utility gain. This estimate will differ from $\bar{d}$ by a Jensen inequality term because of the non-linearity of the function, as shown in Figure 1.

Another issue worth mentioning relates to the omission of the information on financial market participants, which brings in the estimation a potential source of bias due to the censoring of the expected utility gain, $y^*_{h,t}$. If $y^*_{h,t}$ sample mean differs from the population mean simply because the composition of the sample is different, the estimates of the fixed cost bounds based only on data on non-participants will be biased. The issue can be addressed by identifying the selection rule and correcting for the possibility of selection bias by means of an equation explaining *initial* participation such as a latent variable model predicting asset holdings when the portfolio decision takes place[9]. Thus, let r^* be an underlying latent variable denoting the level of indirect utility associated to the portfolio choice of interest:

$$\begin{aligned} r^*_{h,t} &= \kappa_{h,t}'\pi + \eta_{h,t} \\ &= V_{h,t} + \eta_{h,t} \end{aligned} \tag{4}$$

where κ is a *kx1* vector of household specific observable characteristics and $\eta_{h,t}$ is a household specific unobservable variable. For $r^*_{h,t} > 0$, participation occurs, in which case a dichotomous variable, $D_{h,t}$, is equal to one; otherwise, it is zero. Then, if $\eta_{h,t}$ and the individual random taste shifter, $u_{h,t}$, are distributed jointly as standard normal random variables and $E_{(h)}\{\exp(u_{ht}) \mid D_{h,t} \neq 1\} = E_{(h)}\{\exp(u_{ht+1}) \mid D_{h,t} \neq 1\}$[10], the mean value of the expected utility gain in the sub-sample excluding participants can be written as (omitting the superscript *i*):

[9] It is worth pointing out that the expected utility gain equation in (1) does not determine the household type. It simply ensures the non-participants are happy to hold on to their choices.

[10] This implies assuming that $\exp(u_{h,t})$ is a random walk in the sub-sample considered.

$$
\begin{aligned}
E_{(h)}\{y_{h,t} \mid D_{h,t} \neq 1\} &= E_{(h)}\{E_t[\tilde{v}_{h,t+1}(x,\delta)] \cdot \exp(u_{h,t}) \mid D_{h,t} \neq 1\} \\
&= E_t[\tilde{v}_{h,t+1}(x,\delta)] \cdot E_{(h)}\{\exp(u_{h,t}) \mid D_{h,t} \neq 1\} \\
&= E_t[\tilde{v}_{h,t+1}(x,\delta)] \cdot s_{h,t}(V, \rho_{u\eta})
\end{aligned} \tag{5}
$$

where $E_{(h)}\{\ \}$ is the mean taken across households, whereas $E_t[\ ]$ is household conditional expectation. Also,

$$
\tilde{v}_{h,t+1}(x,\delta) = U(\tilde{c}_{h,t}(x,\delta)) - U(c_{h,t}) + \beta[U(\tilde{c}_{h,t+1}(x,\delta)) - U(c_{h,t+1})] \tag{6}
$$

and[11]

$$
s_{h,t}(V, \rho_{u\eta}) = e^{1/2} \frac{1 - \Phi(V_{h,t} + \rho_{u\eta})}{1 - \Phi(V_{h,t})} \tag{7}
$$

where Φ refers to the cumulative standard normal and $\rho_{u\eta}$ is the correlation between $u_{h,t}$ and $\eta_{h,t}$. Thus, the model corrected to account for sample selection can be written as:

$$
y_{h,t} = E[w_{h,t}\tilde{v}_{h,t+1}(x,\delta)] \cdot s_{h,t}(V, \rho_{u\eta}) + \xi_{h,t} \tag{8}
$$

where $\xi_{h,t}$ is an error, such that $E_h\{\xi_{h,t} \mid D_{h,t} \neq 1\} = 0$. In practice, when bounding the cost associated to the market for risky assets, the sample selection correction term will account for the exclusion of risky asset holders; when bounding the costs associated to the market for riskless assets, it will account for the exclusion of those who hold such assets.

3.2. Estimation Procedure

The estimation of the parameters of interest takes two steps. In the first step, I evaluate the sample selection correction term, $s_{h,t}(V, \rho_{u\eta})$. Then, after substituting it in (8), in the second step, I estimate the household optimal portfolio and the transaction cost bound using a method of moment estimator. The sample selection correction term entails the identification of two sets of parameters: the coefficients of the household specific observable characteristics in the latent variable model for portfolio choice, in (4), and the unobservable correlation between $u_{h,t}$ and $\eta_{h,t}$, $\rho_{u\eta}$. The first set of parameters can be obtained by maximum likelihood estimation of the bivariate probit associated to the latent variable model. The unobservable correlation between $u_{h,t}$ and $\eta_{h,t}$ can hardly be identified and distinguished from the unknown

[11] See the Appendix for the derivation of this result.

parameters that enter the expected utility gain function, given the multiplicative structure of the model in (8). However, since $\rho_{u\eta} \in [-1,+1]$, I can proceed and determine the range of values that the cost bounds can take on depending on the value of $\rho_{u\eta}$. Under the assumption of isoelastic utility, another parameter that cannot be identified within the model is the coefficient of relative risk aversion characterizing household preferences. As I do with $\rho_{u\eta}$, I assign relative risk aversion a range of values and verify how sensitive my estimates are to such parameter.

3.2.1. The Optimal Portfolio

In order to identify the potential gains from financial market participation and measure the transaction cost bounds, the household optimal portfolio, x, must be determined. Let $x_{h,t}(g) = f(z_{h,t}, g) \cdot c_{h,t}$, where z is an $mx1$ vector of instruments that have been shown useful in predicting market returns; z varies over time and can be household specific. $f(.)$ is a logit transformation of an $mx1$ vector g of parameters[12]. The household optimal portfolio is simply the investment ensuring the maximum return in terms of utility, given the per-period participation costs. Thus, it can be estimated by maximizing households utility in case of participation with respect to the vector of unknown parameters, g, given the fixed transaction cost, i.e. by solving the following problem:

$$\max_g E\left[U(\tilde{c}^i_{h,t}(x(g),\delta)) \cdot s_{h,t}(V,\rho_{u\eta}) + \beta \cdot U(\tilde{c}^i_{h,t+1}(x(g),\delta)) \cdot s_{h,t}(V,\rho_{u\eta})\right]$$

As I have mentioned before, the optimal portfolio is determined by exploiting the fact that asset returns are to some extent predictable. Since, in practice, the vector of instruments z that I use does not vary across households, but varies only over time, optimal holdings cannot be estimated by exploiting across household variability, but only the variability over time. Thus, I can compute x by solving:

$$\max_g E\left[\overline{U(\tilde{c}^i_{h,t}(x(g),\delta)) \cdot s_{h,t}(V,\rho_{u\eta})} + \beta \cdot \overline{U(\tilde{c}^i_{h,t+1}(x(g),\delta)) \cdot s_{h,t}(V,\rho_{u\eta})}\right] \tag{9}$$

where

$$\overline{U(\tilde{c}^i_{h,t}(x(g),\delta)) \cdot s_{h,t}(V,\rho_{u\eta})} = H_t^{i^{-1}} \sum_{h_t} U(\tilde{c}^i_{h,t}(x(g),\delta)) \cdot s_{h,t}(V,\rho_{u\eta})$$

[12] Specifically: $f(z,g) = (1+\exp(z'g))^{-1}$. This specification is dictated primarily by computational considerations.

is time t mean household utility. H^i_t is the number of households of type i who had their first interview in the tth time period. If the maximand is sufficiently smooth and an optimal portfolio, $x(g)^*$, associated to the fixed cost exists, then, in terms of first-order conditions, the optimal g must be such that (equation (10))

$$(T-1)^{-1}\sum_{t=1}^{T-1}\left[D_g\overline{U(\tilde{c}^i_{h,t}(x(g),\delta))\cdot s_{h,t}(V,\rho_{u\eta})}+\beta\cdot D_g\overline{U(\tilde{c}^i_{h,t+1}(x(g),\delta))\cdot s_{h,t}(V,\rho_{u\eta})}\right]=0$$

where D_g is the derivative with respect to g.

The idea behind the optimal portfolio estimation procedure is that of capturing the unexploited investment opportunities for non-participants using their own mean utility as pricing kernel. Thus, by solving the set of equations[13] in (10) and focussing on those who do not hold risky assets, I can determine their optimal investment in such securities (given the costs) in case of participation to the market. Similarly, by focussing on those who do not hold riskless bills, I can determine their optimal investment is such assets. Notice that, in practice, the actual transaction costs are not observed, nor estimated and only the cost bounds are identified. Therefore, the optimal portfolios of risky and riskless assets are determined as a function of a level of costs equal to the estimated bounds, which is consistent with the rest of the analysis.

3.2.2. The Transaction Cost Bounds

The estimation of the cost bounds is based on the conditional moments in (3), which, after correcting for sample selection, can be written as:

$$E[\tilde{v}^i_{h,t+1}(x,d)\,|\,I_{h,t}]\cdot s^i_{h,t}(V,\rho_{u\eta})+\xi_{h,t}=0 \qquad i=2,3;\quad \forall t \tag{11}$$

where $I_{h,t}$ is household h information set at time t. Let $W_{h,t}$ be a collection of non-negative[14] variables in $I_{h,t}$ observable to the econometrician. Taking any $w_{h,t}$ in $W_{h,t}$, it follows from (11) that

$$E[w_{h,t}\tilde{v}^i_{h,t+1}(x,d)]\cdot s^i_{h,t}(V,\rho_{u\eta})+\xi_{h,t}=0 \qquad i=2,3;\quad \forall t \tag{12}$$

[13] The first-order conditions 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, the second-order condition must be checked as well.

[14] The non-negativity assumption is not strictly needed. However, in order to ensure that the inequality implied by (1) has the same sign across households, it is necessary that the variables in W_{ht} have the same sign across households.

As mentioned in the previous section, the lack of a longer panel dimension in the data set precludes estimating the individual cost bounds, *d*. However, by aggregating properly across households, we can identify the bound to the costs for a consumer whose expected utility gain equals the mean expected gain. Then, the relevant moment conditions are:

$$E\left[E_{(h)}\left\{w_{h,t}\tilde{v}^{i}_{h,t+1}(x,d)\cdot s^{i}_{h,t}(V,\rho_{u\eta})\right\}\right]=0 \qquad i=2,3 \tag{13}$$

which yield a consistent estimator of the bounds if the trading rules as a function of the parameters are well behaved and if $w_{h,t}\tilde{v}^{i}_{h,t+1}(.)$ is time stationary and has finite mean, so that some law of large numbers can be applied. By means of (13), it is possible to estimate consistently the fixed cost lower bound, $d\leq\delta$, as a function of $\rho_{u\eta}$ and of the coefficient of relative risk aversion.

4. Empirical Analysis

4.1. Data

The data used to estimate the fixed cost bounds are taken from the US Consumer Expenditure Survey (CEX), which is a representative sample of the US population, run by the US Bureau of Labor Statistics. The survey is a rotating panel in which each consumer unit is interviewed every three months over a twelve months period, apart from attrition. The data used for the analysis cover the period 1982-1995 and the sample consists of 24,643 households. Each quarterly interview collects household monthly expenditure data on a variety of goods and services for the three months preceding the one when the interview takes place. In the final interview, an annual supplement is used to obtain a financial profile of the household providing information as to the amounts held in checking, brokerage and other accounts, in saving accounts, in US saving bonds and as to the market value of all stocks, bonds, mutual funds and other securities. The changes occurred in such stocks over the previous twelve months are also reported.

The consumption measure I use is deseasonalized, real monthly per-adult equivalent[15] expenditure on non-durable goods and services. Given the timing of the data on asset holdings, for each household only two consumption observations are

[15] 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 in the households are weighted 0.8 and the children are weighted 0.4.

used: the one for the month preceding the first interview and the one for the month preceding the last, implying a nine-month gap. It follows that for each household only a single observation on the expected utility gain, $E_t[v_{h,t+1}(.)]$, can be defined. t is the month of the first observation on consumption and $t+1$ that of the second. Since the interviews take place throughout the year, in the sample used, t runs from 1981:12 to 1985:5 and from 1986:1 to 1994:12, for a total of 150 periods[16].

The household type is determined on the basis of asset holdings twelve months before the last interview, which can be computed by subtracting the changes occurred over that period to the stocks held at the time of the last interview. The variables "stocks, bonds, mutual funds and other securities" and "US saving bonds" are added together and those households who report a non-null amount of such variable are defined as risky asset holders. As a measure of riskless asset holdings, I take the amounts held in checking and saving accounts. Table 1 reports the sample composition in each of the years considered on the ground of household asset portfolios. The first column of the Table contains the share of households holding a positive amount of both risky and riskless assets. They represent about 30.5% of the sample. The second column reports the share holding only riskless assets (51% of the sample). The third column indicates how many households do not hold either asset (18.5% of the sample). 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. This is consistent with the evidence found by Poterba and Samwick (1997) using the US Survey of Consumer Finance, which suggests that equity ownership has increased over time especially through mutual funds and tax-deferred accounts. Also, they find a sharp rise in the fraction of households owning both tax-exempt and taxable bonds.

Table 2 reports some descriptive statistics for the sample as a whole and for the three types of households. Type 1 households, who participate to both markets, 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 monthly family income is higher, as well as their per-capita consumption. Those who hold neither

[16] See the Appendix for an explanation of the discontinuity and for further details on the data, on household selection and exclusions and on variable definition.

risky nor riskless assets tend to be the least educated and to have the lowest income and consumption and in 41 percent of the cases are headed by a woman.

Asset returns are summarized in Table 3. As risky return I take the total return (capital gains plus dividends) on the S&P500 Composite Share Index. As riskless return I take the return on US Treasury bills. The data in the Table are returns over the nine-month period that runs between the two consumption observations used in the analysis. The mean equity premium over the sample period considered is about seven percent.

4.2. Estimation Results

4.2.1. The Sample Selection Correction Term

Before estimating the fixed cost bounds, the sample selection term, $s_{h,t}(V, \rho_{uv})$, must be determined to account for the censoring of the expected utility gain. Such objective can be achieved by means of a bivariate probit model for participation at time *t*. The variables included as determinants of the probability of asset holding are a polynomial in age, a set of education dummies, the education dummies interacted with age, a dummy for the presence of children, a dummy for single person households and a dummy for the region where the household resides. Fourteen year dummies are also included. The first column of Table 4 reports the estimation results for the probability of participating to both financial markets. Such probability appears to increase non-linearly with age and with education; it is higher among male-headed households and is lower among single person households. These estimates allow to construct the sample selection correction term for the case when the analysis is based on those households who do not hold risky securities to bound the costs of participating to the market for such assets. The second column of the Table reports the estimation results for the probability of holding either both assets or no assets at all, which corresponds to one minus the probability of holding only riskless securities. These figures allow to correct for sample selection when computing the risky asset market cost bound using only the information on those households who do not hold risky assets, but do hold some riskless ones. Given the apparent disparities between those who participate to both markets (type 1) and those who do not participate to any (type 3), the results from the estimation are not as clear-cut as those reported in the first column and their interpretation is not as straightforward. The last column of the

Table shows the results for the probability of holding either both assets or only some riskless assets and allows to correct for sample selection when the analysis is based on those who do not hold any financial securities. The outcome is very similar to that reported in the first column, both from a qualitative and a quantitative point of view, with households having an older, more educated and male head more likely to participate to financial markets.

In order to compute the sample selection correction term, as defined in (7), a value must be assigned to the unobservable and non-identifiable correlation between $u_{h,t}$ and $\eta_{h,t}$, which, in the tables below, is set equal to +0.5 and -0.5 to assess the effect of a positive correlation in the first case and of a negative one in the second.

4.2.2. The Optimal Portfolio and The Transaction Cost Bounds

Three sets of results are presented in this section. The first set refers to the costs of participating to the market for risky assets; the second looks at the possible costs of participating to the market for riskless ones and the third set focuses on the two markets considered jointly. Once determined the appropriate sample selection correction term, moment conditions (10) and (13) can be used to estimate jointly the optimal portfolio and the lower bound to the per-period cost of participating to the market of interest. For identification purposes, two sets of instruments are needed. The first set (z), identifying the parameters defining the optimal investment at time t, includes the returns on the S&P500 CI and on Treasury bills, the rate of growth of GDP and the rate of inflation. All variables are lagged one period and refer to the time interval from $t-1$ to t. The second set (w), consisting of good predictors of the utility gains in case of participation, includes household monthly consumption and income at time t, a second order polynomial in the household head age, two education dummies for household head with high school diploma and university degree and all the instruments in z (plus a constant). Thus, the estimation relies on 15 equations to determine 5 parameters, which provide the basis for an overidentifying restriction test of the model.

The Tables 5 to 9, reporting the estimates of the parameters of interest, have the following structure. The results in panel (a) are obtained by setting $\rho_{u\eta}=0.5$, those in panel (b) by setting $\rho_{u\eta}=-0.5$. Each column is computed assuming isoelastic preferences for different levels of risk aversion. The first row of each table reports the estimates of the bound to the fixed per-period participation costs in dollars of year

2000. These are annualized figures obtained by multiplying by twelve the GMM estimates that are based on monthly consumption data and, therefore, are an average of the mean monthly utility gain over the sample period considered. The reason for multiplying these estimates by twelve is to relate the gains from financial market participation to annual expenditure. The next set of rows in the tables contains the estimates of the parameters of the optimal asset portfolio, which implies investing in the financial market the share of time t consumption reported in the row before the last. The shares reported are average values; in fact, the portfolio parameters are determined using time-varying instruments and consequently the optimal shares to invest vary over time. Standard errors are reported in parentheses[17]. The Sargan test of overidentifying restrictions is reported in the last row. The rate of discount over the nine-month period of investment, β, is set equal to 0.98, which implies an annual rate of approximately 0.97. A nine-month rate of 0.99 implies slightly higher bound estimates, but the overall conclusions do not change in any significant way.

a. The optimal portfolio

Table 5 and 6 focus on the market for risky assets. The results in Table 5 are obtained by focussing on all households who do not report holdings of risky assets; those in Table 6 are obtained using only the information on those who do not hold risky assets, but still hold some riskless securities. The figures reported in the two tables are very similar. Notice also that there are almost no differences between the two panels of the tables, suggesting little sensitivity to the value assigned to $\rho_{u\eta}$, which is chosen arbitrarily. According to the estimates in the first column of the tables, a household with a relative risk aversion of 0.5 could benefit from participating to the risky asset market and optimal behavior would involve investing around 12.5% of current consumption. The literature on portfolio choice predicts that as risk aversion increases households should reduce their risky asset investments. This is unequivocally supported by the evidence displayed in the Tables, according to which as the coefficient of risk aversion increases, the optimal portfolio as a share of consumption falls rapidly. If risk aversion is 3.5, the optimal portfolio of risky asset

[17] The standard errors have been corrected to account for the MA(9) structure of the error due to the overlapping of the observations on the expected utility gain, which follows from the monthly frequency of interviews. The issue of non-positive definite variance covariance matrix in finite samples has been taken care of by using a set of weights like in Newey and West (1987).

should correspond to just around 5% of consumption. The standard errors of the portfolio parameters reported in the Tables suggest that the coefficients associated to the instruments are generally statistically significant.

Table 7 reports the results of an exercise aimed at quantifying the downward bias in the transaction cost estimates reported in Table 5. As mentioned in Section 2, the gains from participation are likely to be under-estimated because of the unavailability of a measure of cash-in-hand. Given the information on total after-tax family income, it is possible to make an assumption as to the wealth held in liquid lower-return assets that is likely to be either invested in the risky asset or consumed immediately, once paid the participation cost and gained access to the higher return risky asset market[18]. The estimates reported in the Table result from the assumption that households can reallocate one percent of monthly income and that the savings they reallocate are initially invested in a zero return asset. The one percent income figure is low; yet, it seems reasonable since total after-tax income does not account for mortgage payments, health insurance, retirement contributions, etc. which limit considerably the amount of liquid wealth immediately available for reinvestment. Also, one percent of income corresponds to approximately 4.5 percent of the monthly consumption figures used in the analysis.

The Table reports the optimal portfolio as a share of "estimated" cash-in-hand: a household with a relative risk aversion of 0.5 should invest in the risky asset market around 16 percent of its cash-in-hand; one with a risk aversion of 3.5 should invest around 10 percent. Compared to the figures in Table 5, those in Table 7 suggest that, if households can invest in the risky asset market also part of their accumulated wealth, they will reduce their consumption slightly less if they are little averse to risk and relatively more if they are more risk averse. Yet, overall, the differences in terms of reallocation of current consumption between the two sets of Tables are rather small - less than one percentage point – and are the result of the complex interaction of the investment riskiness with the fact that transaction costs are fixed and more resources are now available for investment.

To verify whether there are differences in the costs of participating to different financial markets and to get some sense of the magnitude of these differences, I have used the model for the expected utility gain to determine the benefit that those

[18] See footnote 8 in Section 2.

households who hold neither risky nor riskless portfolios would reap by investing in riskless assets. The set of results, shown in Table 8, is obtained by focussing on these agents and using moment conditions (10) and (13) to estimate jointly the optimal portfolio and the lower bound to the costs of participating to such market. The portfolio parameters estimates suggest that a household with a risk aversion of 0.5 could increase its utility by investing in the riskless asset market around 8 percent of its consumption. As before, as risk aversion increases, the utility maximizing investment decreases, but the rate of decrease is much lower than in the case of a risky asset portfolio. The standard errors of the coefficients associated to lagged returns are generally statistically significant, whereas the evidence on the coefficients of GDP and inflation is mixed, suggesting that the latter have little additional predictive power over lagged asset returns.

Finally, I have considered the case where households holding neither risky nor riskless portfolios are allowed to invest in both (or either) assets after paying a fixed cost unrelated to the specificity of the investment. As instruments to determine the optimal share of consumption to invest in financial assets, I use lagged returns on T-bills and on the S&P500 CI and lagged GDP growth and inflation. To compute the optimal portfolio share of risky assets, I use the equity premium lagged one period[19].

Table 9 reports the results of the estimation. As before, as risk aversion increases, the total optimal investment in risky and riskless assets as a share of consumption decreases. Yet, the portfolio parameter estimates exhibit an interesting feature: in fact, they suggest that if the costs were low enough, households would choose to participate to financial markets by holding an optimal portfolio consisting almost exclusively of risky assets. Only, for a coefficient of risk aversion equal to 2.5 or higher, riskless asset holdings become non-negligible. This result suggests that, if they can choose between risky and riskless assets, households clearly prefer the former, which could be expected given the assumption of single fixed participation cost *vis á vis* the considerable risk premium. Yet, as risk aversion increases, the high volatility of risky returns makes these securities less desirable and households rapidly reduce their risky asset holding. At the same time, they start investing in riskless

[19] For computational reason, both the overall investment as a share of consumption and the share invested in risky assets are determined by means of logit transformations of the coefficients of the instruments (see footnote 12). This implies ruling out the possibility of borrowing at the riskless rate and saving in the risky asset.

assets which provide a convenient means of smoothing consumption at a very low risk. As it will be discussed more thoroughly later, the benefits related to consumption smoothing can be expected to be rather important for this group of households, whose expenditure at time *t+1* is lower than that at *t* by 10% on average.

Before turning to the results concerning the lower bounds to the costs of participating to financial markets, it is worth addressing the issue of the sensitivity of all the estimation results to the value[20] taken on by $\rho_{u\eta}$ in the sample selection correction term. Under self-selection, those individuals who have a "comparative disadvantage" with financial market participation will not hold financial assets because their gain is lower than that of a randomly selected sample of individuals with the same characteristics. Thus, the need to control for the exclusion of asset holders when estimating the participation gain. The lack of sensitivity to the value taken on by the correlation between the unobservable of the model for the utility gain from participation and the unobservable affecting the likelihood of participating can be interpreted as evidence of very limited self-selection.

b. The transaction cost bounds

The discussion in the previous section on optimal portfolios showed that those who have chosen not to hold one or more of the available securities could increase their utility by participating optimally to the relevant markets. Yet, if participation costs are high enough, any gain would be eliminated and non-participation becomes optimal.

Table 5 and 6 report the estimates of the lower bound to the costs rationalizing non-participation to the risky asset market. According to the figures in Table 5, panel (a), a household with relative risk aversion of 0.5 would not net any positive gain from participating optimally to the risky asset market if the annual costs involved were greater than $91. As risk aversion increases, the estimated bounds decrease at a falling rate and tend to level off for coefficients of risk aversion above 2.5/3. This trend in the estimates results from the fact that the lower bound is a measure of the gains from participation and, when the investment is risky, such gains are lower the more concave the utility function. The standard errors reported in parenthesis imply

[20] In addition to ±0.5, other values (not reported for brevity) have been tried. No important difference in either the portfolio parameters or the bounds has been recorded.

that the bounds are estimated with acceptable precision. In fact, the 95% confidence intervals range from approximately $45 to $137 for a risk aversion of 0.5 and from $53 to $80 for a risk aversion of 3. Finally, the Sargan test, whose p-value is reported in parenthesis in the last row of the Table, never rejects the overidentifying restrictions to the model, providing evidence in favor of the hypothesis of correct specification. Notice, that, like with portfolio parameters, there are negligible differences between Table 5 and 6 and also between the two panels of each Table.

As pointed out earlier, these figures are obtained without accounting for household cash-in-hand and for the possibility of reinvesting accumulated savings. As a consequence, they can be expected to be somewhat downwardly biased. Table 7 addresses the issue of the severity of this problem. According to the Table, a household with a risk aversion of 0.5 that can reinvest one percent of its after-tax income, in addition to reallocating its consumption expenditure, will not net a positive gain from optimal risky asset holding if the annual participation costs are higher than approximately $148. The set of estimates of the gains from participation in Table 7 is to some extent higher than those seen before, as it could be expected given the fixed nature of the costs *vis-á-vis* the fact that now households have more resources to invest. Yet, they remain reasonably low to be thought to bound actual market frictions. Also, they can be expected to fall rapidly when assuming that the accumulated saving that are to be moved into the costly security were invested in a positive return asset instead of a zero return one, like I have assumed here. According to the figures in Table 7, as risk aversion increases, the estimated bound does not change significantly.

Next, to address the issue of the differences in the costs of different financial portfolios, I have estimated the gains from participating to the riskless asset market, using the information on those households who hold no financial securities. The point estimates of the bounds, reported in Table 8, are always strongly significant, suggesting that investing in riskless assets is also somewhat costly. According to the Table, for a household with a risk aversion of 0.5, it is optimal not to participate to the riskless asset market, if participation involves costs that are higher than around $24. These figures suggest that the gains from holding riskless assets are quite small and, as expected, they are significantly smaller than those recorded for risky asset market participation, at least for low levels of risk aversion. Yet, they tend to increase rapidly as risk aversion increases: in fact, for a household with a risk aversion of 3.5 the

bound estimate is above $63, which is of the same order as the bound for non-participants to the risky asset market with similar risk aversion (see Table 6). As to the precision of the bound estimates as measured by the width of the confidence intervals, like in the previous tables it appears to be negatively correlated to the size of the bound. However, in the case of riskless asset markets, it appears to be slightly larger, with somewhat narrower confidence intervals. Finally, as before, a Sargan test of the over-identifying restrictions never rejects the null of correct specification of the analysis.

The positive relationship between the bound estimates and the degree of risk aversion is due to the specific nature of the gains from having access to a riskless security. In fact, the main benefit in terms of utility from investing in such assets comes from the possibility of smoothing consumption over time, without increasing significantly consumption risk, although life-time consumption does not rise significantly because of the limited size of the returns. The more risk averse the agent, the greater the utility gain from smoothing consumption, the higher the bound to the cost rationalizing non-participation. As I have mentioned before, in the sample I use those who do not hold riskless assets exhibit falling consumption, on average. Such behavior can hardly be rationalized within the standard neo-classical model for consumption, according to which these households would undoubtedly benefit from smoothing consumption by investing in a riskless asset. Yet, if the costs involved in shifting consumption over time are higher than the estimated bound, their choices can be fully rationalized.

The last type of analysis I have carried out aims at determining the gain from having access to a market where both risky and riskless securities can be traded. The gain represents the lower bound to the single fixed cost rationalizing the behavior of those households in the sample that have chosen not to hold any financial asset. Table 9 reports the results of the estimation: for a household with 0.5 relative risk aversion, the point estimate of the bound is approximately $60. As risk aversion increases, the lower bound at first does not change or decreases somewhat, but then start increasing and for a risk aversion coefficient of 3.5, it is around $75. Overall, the results in Table 9 are consistent with those in the previous tables and shed further light on the nature of the gains and, therefore, on the lower bound to the costs of financial market participation. In fact, the trend in the bound estimates, together with those in the portfolio parameters suggest that the nature of the gain is different at different levels

of risk aversion. As discussed in the previous section, if the participation cost is unrelated to the type of investment and households can choose between risky and riskless assets, they appear to prefer the former, which could be expected given the considerable risk premium. Yet, as risk aversion increases, the utility benefit from holding a risky portfolio for its high expected return falls rapidly and households tend to reduce their investment because of the high volatility of the risky returns. Besides rising expected life-time consumption, risky assets can also provide a means of smoothing consumption, which is particularly valuable for the sample of households considered here and this helps explain the increase in the bound that can be recorded for values of risk aversion above 2/2.5. Also, at these levels of risk aversion, which is when consumption-smoothing considerations appear to become important, households do not just reduce their overall investment, but also switch to riskless assets, which provide a convenient means of smoothing consumption at a very low risk. Finally, notice that for levels of risk aversion above 2, the bound estimates in Table 9 are comparable to those in Table 8, which refer to the market for riskless assets, although they are somewhat higher, probably as a result of the higher return of the means available to smooth consumption.

5. Concluding Remarks

This paper focuses on the issue of limited participation to financial markets and determines a lower bound on the level of fixed participation costs that is required to reconcile observed consumption choices with asset returns within an isoelastic utility framework. The bound is obtained from the set of conditions that ensure the optimality of observed behavior for financial market non-participants. The evidence found suggests that reasonably low costs can justify observed behavior for degrees of risk aversion held as realistic by the literature. In fact, under the assumption of log utility, conservative estimates corresponding to the upper extreme of 95 percent confidence intervals imply a lower bound to the annual fixed costs that rationalize non-participation to risky assets markets in the range of 95-175 dollars, which amounts to less than 0.4-0.7% of household annual consumption. To justify non-participation to riskless asset markets, smaller frictions are sufficient.

An interesting point that has emerged from the analysis is that for many households most of the gains from financial market participation are not as much related to the size of the returns, as to the benefits from smoothing consumption.

However, overall, for the sample of non-participants considered here, the gains from participating to financial markets do not appear to be large enough to justify the investment *vis-á-vis* the costs. The results based on a "guess" of household cash-in-hand suggest that this might be due to the fact that the resources available for investment are limited. Yet, the differences in terms of wealth between participants and non-participants do not seem wide enough to justify such different asset holding behavior. A more reasonable explanation can instead be found in the amount of household heterogeneity, both observable and unobservable, which appears to explain the differences in the consumption patterns across household types. In fact, participants and non-participants are likely to differ in terms of taste for risk, individual ability, faculty of modifying labor supply, etc. Differences in all these factors can be expected to have an effect both on the gains from asset holdings and on the costs of financial market participation, in which case, the kind and the size of the benefits of observed financial market participants will be very different and much larger than those recorded for non-participants, whereas the level of costs, especially of figurative charges, can be expected to be much smaller.

References

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

Attanasio, O., Banks, J. and S. Tanner, (1998), Asset Holding and Consumption Volatility, IFS Working Paper W98/8.

Blundell, R., Reed, H. and T. Stocker, (1998), The Evolution of Aggregate Wages: Heterogeneity, Selection and Distribution, mimeo.

Cragg, J., (1971), Some Statistical Models for Limited Dependent Variables with Application to the Demand for Durable Goods, *Econometrica*, vol. 39, 829-44.

Davidson, R. and J. MacKinnon, (1993), *Estimation and Inference in Econometrics*, Oxford University Press.

Haliassos, M. and C. Bertaut, (1995), Why Do So Few Hold Stocks?, *Economic Journal*, vol. 105(432): 1110-1129.

Hansen, L. and R. Jagannathan, (1991), Implication of Security Market Data for Models of Dynamic Economics, *Journal of Political Economy*, vol. 99: 225-262.

Hansen, L. and K. Singleton, (1983), Stochastic Consumption, Risk Aversion and the Temporal Behavior of Asset Returns, *Journal of Political Economy*, vol. 91: 249-265.

Hausman, J. and D. Wise, (1978), A Conditional Probit Model for Qualitative Choice: Discrete Decisions Recognizing Interdependence and Heterogeneous Preferences, *Econometrica*, vol. 46, 403-26.

Hausman, J. and D. Wise, (1979), Attrition Bias in Experimental and Panel Data: The Gary Income Maintenance Experiment, *Econometrica*, vol. 47, 455-73.

Hsiao, C., (1986), *Analysis of Panel Data*, Cambridge University Press.

Keane, M., Moffit, R. and D. Runkle, (1988), Real Wages over the Business Cycle: Estimating the Impact of Heterogeneity with Micro Data, *Journal of Political Economy,* vol. 96, 1232-66.

Kocherlakota, N., (1990), The Equity Premium: It Is Still a Puzzle, *Journal of Economic Literature*, 24, 42-71.

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

Maddala, G. (1983), *Limited-Dependent and Qualitative Variables in Econometrics*, Cambridge University Press.

Mankiw, N. and S. Zeldes, (1991), The Consumption of Stockholders and Non-Stockholders, *Journal of Financial Economics*, 17, 211-19.

Mehra, R. and E. Prescott, (1985), The Equity Premium Puzzle, *Journal of Monetary Economics*, vol. 15: 145-61.

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

Nijman, T. and M. Verbeek, (1992), Nonresponse in Panel Data: The Impact on Estimates of a Life Cycle Consumption Function, *Journal of Applied Econometrics*, vol. 7: 243-57.

Nijman, T. and M. Verbeek, (1996), Incomplete Panels and Selection Bias, in Matyas, L. and P. Sevestre (eds.), Kluwer Academic Publishers, (2nd ed.): 449-90.

Paiella, M., (1999), Partecipazione ai Mercati Finanziari e Volatilitá dei Consumi: Una Soluzione all'Equity Premium Puzzle?, mimeo, University of Pavia, Italy.

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

Vissing- Jørgensen, A., (1999), Limited Stock Market Participation and the Equity Premium Puzzle, mimeo, University of Chicago.

Vissing-Jørgensen, A., (1999), Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures, mimeo, University of Chicago.

Appendix

Data Description

The data used to estimate the fixed cost bounds are taken from the US Consumer Expenditure Survey (CEX), which is a representative sample of the US population, run by the US Bureau of Labor Statistics. The survey is a rotating panel in which interviews take place throughout the year and each consumer unit is interviewed every three months over a twelve months period. This rotating procedure is designed to improve the overall efficiency of the survey and to reduce the problems of attrition. New households are introduced into the panel on a regular basis as old ones complete their participation and, as a whole, about 4500 households are interviewed each quarter, more or less evenly spread over the three months.

The data used for the analysis cover the period 1982-1995. I exclude from the sample those households with incomplete income responses and those living in rural areas or in university housing. In addition, I exclude those whose head was less than twenty-five or older than sixty-five (about 10,000 households), those who do not participate to all interviews (about 33% of the initial sample), the top 0.1 percent and the bottom 1.7 percent of the income distribution. The reason for this latter selection is to exclude about 500 households who report a total after-tax income below $3,500 and who are likely to consume all their income and have no resources to invest in financial markets. Finally, I select out those households with average monthly per-adult equivalent consumption[21] 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 over the nine months period, $\Delta c_{h,t}$, greater than $1,750 in absolute value (about 500 households). For several households the financial supplement contains many invalid blanks either in the stocks of assets or in the changes occurred with respect to the previous year. Since I am interested in the asset holding choice, - and not in the actual amounts held -, I keep not only those households who report both a "valid stock" and a "valid change", but also those who report only one of the two amounts of interest[22]. Overall, the sample used consists of 24,643 households.

[21] See footnote 15 in Section 4.

[22] About 3,000 households report invalid information in either the flows or stocks of financial assets.

The consumption measure I use is deseasonalized, real[23] monthly per-adult equivalent expenditure on non-durable goods and services. The exclusion of durable consumption is grounded in the assumption of separability of preferences between durables and non-durables/services. Given the timing of the data on asset holdings, for each household only two consumption observations are used: the one for the month preceding the first interview and the one for the month preceding the last, implying a nine-month gap. It follows that for each household only a single observation for the expected utility gain, $E_t[v_{h,t+1}(.)]$, can be defined. t is the month of the first observation on consumption and $t+1$ that of the second. Because of this matching of households forward in time, a problem arises around 1985-86 when the sample design and the household identification numbers were changed. As a consequence, it is not possible to match forward those households who have their first interview in the third and fourth quarter of 1985 and they are excluded from the sample. Thus, the sample used consists of households who have their first interview between 1982:1 and 1985:6[24] and between 1986:1 and 1995:1, which implies that t, the month of the first observation on consumption, runs from 1981:12 to 1985:5 and from 1986:1 to 1994:12, for a total of 150 periods.

The household type is determined on the basis of asset holdings twelve months before the final interview, which can be determined by subtracting the changes occurred over that period to the stocks held at the time of the last interview[25]. The variables "stocks, bonds, mutual funds and other securities" and "US saving bonds" are added together and those households who report a non-null amount of such variable are defined as risky asset holders. As a measure of riskless assets, I take the amounts held in checking and saving accounts. Less than 0.4% of the households reports only holdings of risky assets: these households are dropped from the sample.

[23] 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 indices 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.

[24] For the first quarter of 1986, the Bureau of Labor Statistics created two files: one based on the original sample design and one based on the new design. After the first quarter, no track is kept of the households in the old sample. Thus, I can match forward only those households in the original sample who had their first interview before July 1985.

Derivation of the Sample Selection Correction Term

By assumption ($u_{h,t}$ $\eta_{h,t}$) is a joint normal random variable: namely

$$\begin{pmatrix} u_{h,t} \\ \eta_{h,t} \end{pmatrix} \sim \mathrm{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \sigma_{u\eta} \\ \sigma_{u\eta} & 1 \end{pmatrix} \right) \tag{A1}$$

Normality implies that

$$u_{h,t} = \sigma_{u\eta}\eta_{h,t} + \zeta_{h,t} \tag{A2}$$

where $\zeta_{h,t}$ is an error term, normally distributed with zero mean and standard deviation σ_{ζ} and orthogonal to $\eta_{h,t}$ by construction. In order to determine the sample selection correction term, $s_{h,t}(V, \rho_{u\eta})$, I must compute the conditional expectation $E\{\exp(u_{h,t}) \mid D_{h,t} \neq 1\}$, where $D_{h,t} \neq 1$ if $\eta_{h,t} \leq -V_{h,t}$. Using (A2) and the fact that $\zeta_{h,t}$ is independent of $\eta_{h,t}$, I can rewrite such conditional expectation as follows

$$\begin{aligned} E\{\exp(u_{h,t}) \mid D_{h,t} \neq 1\} &= E\{\exp(\sigma_{u\eta}\eta_{h,t} + \zeta_{h,t}) \mid D_{h,t} \neq 1\} \\ &= E\{\exp(\sigma_{u\eta}\eta_{h,t}) \cdot \exp(\zeta_{h,t}) \mid D_{h,t} \neq 1\} \\ &= E\{\exp(\sigma_{u\eta}\eta_{h,t}) \mid D_{h,t} \neq 1\} \cdot E\{\exp(\zeta_{h,t})\} \end{aligned} \tag{A3}$$

The last term of (A3) is simply

$$\begin{aligned} E\{\exp(\zeta_{h,t})\} &= \exp\left(\frac{\sigma_{\zeta}^2}{2}\right) \\ &= \exp\left(\frac{1-\sigma_{u\eta}^2}{2}\right) \end{aligned} \tag{A4}$$

where the second equality follows from (A2). The other term of (A3) can be developed as

$$\begin{aligned} E\{\exp(\sigma_{u\eta}\eta_{h,t}) \mid D_{h,t} \neq 1\} &= \int_{-\infty}^{+\infty} \exp(\sigma_{u\eta}\eta_{h,t}) \cdot \phi(\eta_{h,t} \mid D_{h,t} \neq 1) d\eta_{h,t} \\ &= \frac{\int_{-\infty}^{-V_{h,t}} \exp(\sigma_{u\eta}\eta_{h,t}) \cdot \phi(\eta_{h,t}) d\eta_{h,t}}{1 - \Phi(V_{h,t})} \end{aligned} \tag{A5}$$

which follows from the fact that $D_{h,t} \neq 1$ if $\eta_{h,t} \leq -V_{h,t}$. The integral at the numerator can be rewritten as:

[25] When either of the two variables is missing, I define the household as asset holder either if they hold a positive amount of the asset at the time of the fifth interview or if they report a non-null change with respect to the previous twelve months.

$$\int_{-\infty}^{-V_{h,t}} \exp(\sigma_{u\eta}\eta_{h,t}) \cdot \phi(\eta_{h,t}) d\eta_{h,t} = \int_{-\infty}^{-V_{h,t}} \exp(\sigma_{u\eta}\eta_{h,t}) \cdot (2\pi)^{-0.5} \exp\left(-\frac{\eta_{h,t}^2}{2}\right) d\eta_{h,t}$$
$$= \int_{-\infty}^{-V_{h,t}} (2\pi)^{-0.5} \exp\left(\sigma_{u\eta}\eta_{h,t} - \frac{\eta_{h,t}^2}{2}\right) d\eta_{h,t} \tag{A6}$$

Completing the square of the term in the exponent,

$$\sigma_{u\eta} v_{h,t} - \frac{\eta_{h,t}{}^2}{2} = -\frac{1}{2}\left(\eta_{h,t}{}^2 - 2\sigma_{u\eta}\eta_{h,t}\right)$$
$$= -\frac{1}{2}\left(\eta_{h,t} - \sigma_{u\eta}\right)^2 + \frac{\sigma_{u\eta}^2}{2} \tag{A7}$$

which implies that (A6) can be written as ((A8))

$$\int_{-\infty}^{-V_{h,t}} \exp(\sigma_{u\eta}\eta_{h,t}) \cdot \phi(\sigma_{u\eta}\eta_{h,t}) d\eta_{h,t} = \exp\left(\frac{\sigma_{u\eta}^2}{2}\right) \cdot \int_{-\infty}^{-V_{h,t}} (2\pi)^{-0.5} \exp\left(-\frac{1}{2}\left(\eta_{h,t} - \sigma_{u\eta}\right)^2\right) d\eta_{h,t}$$
$$= \exp\left(\frac{\sigma_{u\eta}^2}{2}\right) \cdot \int_{-\infty}^{-V_{h,t}-\sigma_{u\eta}} (2\pi)^{-0.5} \exp\left(-\frac{1}{2}\eta_{h,t}^{2\,*}\right) d\eta_{h,t}^*$$
$$= \exp\left(\frac{\sigma_{u\eta}^2}{2}\right) \cdot \Phi\left(-V_{h,t} - \sigma_{u\eta}\right)$$
$$= \exp\left(\frac{\sigma_{u\eta}^2}{2}\right) \cdot \left[1 - \Phi\left(V_{h,t} + \sigma_{u\eta}\right)\right]$$

Substituting (A8) in (A5) and (A4) and (A5) in (A3), I obtain

$$E\left\{\exp(u_{h,t}) \mid D_{h,t} \neq 1\right\} = \frac{\exp\left(\frac{\sigma_{u\eta}^2}{2}\right)\left[1 - \Phi\left(V_{h,t} + \sigma_{u\eta}\right)\right]}{1 - \Phi\left(V_{h,t}\right)} \cdot \exp\left(\frac{1 - \sigma_{u\eta}^2}{2}\right)$$
$$= \frac{1 - \Phi\left(V_{h,t} + \sigma_{u\eta}\right)}{1 - \Phi\left(V_{h,t}\right)} \cdot \exp\left(\frac{1}{2}\right) \tag{A9}$$

Since both $u_{h,t}$, and $\eta_{h,t}$ have unit variance, I can rewrite (A9) in terms of correlation between $u_{h,t}$ and $\eta_{h,t}$, instead of covariance, and obtain the sample selection term used in Section 2:

$$E\left\{\exp(u_{h,t}) \mid D_{h,t} \neq 1\right\} = s_{h,t}(V, \rho_{u\eta})$$
$$= \exp\left(\frac{1}{2}\right) \cdot \frac{1 - \Phi(V_{h,t} + \rho_{u\eta})}{1 - \Phi(V_{h,t})} \,. \tag{A10}$$

Table 1: Sample Composition

Year	Type 1 (%)	Type 2 (%)	Type 3 (%)	Total households
1982	26.6	55.3	18.1	1957
1983	27.1	54.9	18.0	2004
1984	27.6	54.9	17.5	1987
1985	25.4	54.5	20.1	967
1986	30.2	52.6	17.2	1935
1987	31.2	51.3	17.5	1940
1988	30.8	50.3	18.9	2003
1989	30.3	50.9	18.8	2001
1990	29.3	52.1	18.6	1978
1991	34.3	45.9	19.8	2027
1992	34.7	46.5	18.8	1841
1993	32.8	47.9	19.3	1910
1994	33.7	47.4	17.9	1939
1995	38.3	40.9	20.8	154
Total	30.5	51.0	18.5	24643

NOTE: The relatively small number of households in 1985 is due to the fact that in 1986 the sample design and the household identification numbers were changed. Those households who entered the survey after June 1985 were dropped (or had their identifier changed) before reaching the last interview (see Appendix).

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

	Type 1	Type 2	Type 3	Whole Sample
Age (mean)	43.88	42.32	43.65	43.04
Less than high school (%)	5.79	14.84	32.21	15.29
High school diploma (%)	51.30	59.30	53.41	55.76
College degree (%)	42.91	25.86	14.38	28.95
Male (%)	77.81	69.24	58.96	69.958
Single person (%)	14.31	20.20	18.02	18.00
Married (%)	76.70	63.75	53.56	65.82
Children (%)	49.30	47.89	53.52	49.36
North-East (%)	21.26	18.16	28.09	20.94
Mid-West (%)	27.45	26.12	24.67	26.26
South (%)	26.69	28.36	29.80	28.12
West (%)	24.60	27.35	17.45	24.68
After tax monthly income	$5,499	$3,957	$3,088	$4,267
c_1 (mean)	$925	$774	$697	$797
c_2 (mean)	$927	$770	$631	$789
No. of Observations	7,527	12,555	4,561	24,643

NOTE: All figures are in dollars of year 2000.

Table 3: Average Nine-Month Returns (1981:12-1995:09)

	Mean	Standard Deviation	Min	Max
S&P500CI	0.1208	0.1277	-0.1839	0.5932
T-Bills	0.0488	0.0167	0.0217	0.0886

Table 4: Results of Probit Estimation

	Probit for probability of holding risky assets	Probit for probability of holding either risky and riskless assets or no assets	Probit for probability of holding some asset (risky or riskless)
Age	6.05	1.74	3.67
	(1.82)	(1.66)	(1.96)
Age^2	-4.41	-1.85	-2.39
	(1.76)	(1.61)	(1.90)
Age^3	1.20	0.61	0.53
	(0.55)	(0.50)	(0.59)
High School Diploma	1.23	-0.56	1.22
	(0.14)	(0.10)	(0.10)
College Degree	1.78	-0.57	2.10
	(0.15)	(0.11)	(0.13)
Sex	0.27	0.01	0.31
	(0.02)	(0.02)	(0.02)
Single	-0.22	-0.14	-0.08
	(0.03)	(0.02)	(0.03)
North East	0.04	0.31	-0.45
	(0.03)	(0.02)	(0.03)
Mid-West	0.09	0.16	-0.16
	(0.02)	(0.02)	(0.03)
South	-0.01	0.13	-0.23
	(0.02)	(0.02)	(0.03)
Children	-0.02	0.10	-0.17
	(0.02)	(0.02)	(0.03)
Age*High School Diploma	-0.46	0.42	-0.50
	(0.12)	(0.08)	(0.09)
Age*College Degree	-0.55	0.68	-1.01
	(0.12)	(0.10)	(0.11)
Constant	-4.26	-0.57	-1.50
	(0.62)	(0.55)	(0.65)
p-value Year Dummies	0.0000	0.0000	0.0000
No. of Observations	24,643	24,643	24,643
Pseudo R^2	0.0715	0.0206	0.0851

NOTE: Standard errors in parentheses.

Table 5: Estimates of the Lower Bounds to the Transaction Costs for the Market for the Risky Asset and of the Corresponding Optimal Portfolios (17,116 households)

Panel (a): ($\rho_{u\eta}$=0.5)

		γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound		90.88	71.00	65.71	64.00	64.30	66.54	70.96
		(23.43)	(13.18)	(9.38)	(7.42)	(6.58)	(7.04)	(9.35)
Portfolio parameters	$R_{t-1,t}$	1.57	1.08	0.85	0.70	0.59	0.48	0.39
		(0.82)	(0.52)	(0.42)	(0.39)	(0.40)	(0.43)	(0.50)
	$R^f_{t-1,t}$	0.89	1.62	1.95	2.15	2.28	2.40	2.46
		(0.81)	(0.52)	(0.42)	(0.40)	(0.42)	(0.47)	(0.54)
	$g_{t-1,t}$	-0.09	-0.06	-0.04	-0.03	-0.02	-0.01	-0.00
		(0.03)	(0.02)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
	$\pi_{t-1,t}$	-0.14	-0.10	-0.08	-0.07	-0.05	-0.04	-0.03
		(0.04)	(0.03)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)
Optimal portf. as % of consumption		12.55	8.43	6.94	6.13	5.62	5.27	5.03
Sargan test (*dof = 10*)		13.15	12.78	12.35	11.37	10.07	11.47	11.21
		(0.22)	(0.24)	(0.26)	(0.33)	(0.43)	(0.32)	(0.34)

Panel (b): ($\rho_{u\eta}$=-0.5)

		γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound		89.15	69.88	64.70	63.00	63.37	65.80	70.96
		(22.47)	(12.74)	(9.13)	(7.28)	(6.49)	(6.86)	(8.89)
Portfolio parameters	$R_{t-1,t}$	1.57	1.08	0.85	0.71	0.59	0.48	0.40
		(0.82)	(0.53)	(0.43)	(0.39)	(0.39)	(0.42)	(0.48)
	$R^f_{t-1,t}$	0.90	1.63	1.95	2.15	2.29	2.40	2.47
		(0.82)	(0.53)	(0.43)	(0.39)	(0.41)	(0.45)	(0.51)
	$g_{t-1,t}$	-0.09	-0.06	-0.04	-0.03	-0.02	-0.01	-0.00
		(0.03)	(0.02)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
	$\pi_{t-1,t}$	-0.14	-0.10	-0.08	-0.07	-0.05	-0.04	-0.03
		(0.04)	(0.03)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)
Optimal portf. as % of Consumption		12.47	8.35	6.86	6.06	5.55	5.21	4.99
Sargan test (*dof = 10*)		13.24	12.87	12.54	11.63	9.98	11.68	11.63
		(0.21)	(0.23)	(0.25)	(0.31)	(0.44)	(0.31)	(0.31)

NOTE: The results in panel (a) are obtained by setting $\rho_{u\eta}=0.5$, those in panel (b) by setting $\rho_{u\eta}=-0.5$. Each column is computed assuming isoelastic preferences for different levels of risk aversion and assuming an annualized discount rate equal to 0.97. The first row of each table reports the estimates of the bound to the fixed annualized participation costs in dollars of year 2000. The next set of rows in the tables contains the estimates of the parameters of the optimal asset portfolio, which implies investing in the financial market the share of time *t* consumption reported in the row before the last. The shares reported are average values. Standard errors in parentheses. The Sargan test of overidentifying restrictions is reported in the last row, with p-values based on 10 degrees of freedom in parentheses.

Table 6: Estimates of the Lower Bounds to the Transaction Costs for the Market for the Risky Asset and of the Corresponding Optimal Portfolios (12,555 households)

Panel (a): ($\rho_{u\eta}$=0.5)

		γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound		95.72	73.52	66.62	63.17	61.30	60.92	61.62
		(26.00)	(15.19)	(11.02)	(9.10)	(7.96)	(7.52)	(7.89)
Portfolio parameters	$R_{t-1,t}$	1.50	0.96	0.70	0.55	0.42	0.31	0.20
		(0.84)	(0.53)	(0.42)	(0.37)	(0.35)	(0.36)	(0.40)
	$R^f_{t-1,t}$	1.05	1.87	2.26	2.50	2.69	2.86	3.02
		(0.82)	(0.52)	(0.40)	(0.36)	(0.36)	(0.40)	(0.47)
	$g_{t-1,t}$	-0.10	-0.07	-0.06	-0.05	-0.04	-0.04	-0.03
		(0.03)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)	(0.03)
	$\pi_{t-1,t}$	-0.15	-0.12	-0.11	-0.10	-0.09	-0.09	-0.09
		(0.05)	(0.03)	(0.03)	(0.02)	(0.02)	(0.03)	(0.03)
Optimal portf. as % of Consumption		12.42	8.24	6.71	5.86	5.31	4.92	4.64
Sargan test (*dof = 10*)		12.48	12.66	13.06	13.01	12.48	11.91	11.30
		(0.25)	(0.24)	(0.22)	(0.22)	(0.25)	(0.29)	(0.33)

Panel (b): ($\rho_{u\eta}$=-0.5)

		γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound		91.91	71.12	64.78	61.69	60.08	59.60	60.22
		(24.25)	(14.23)	(10.53)	(8.54)	(7.56)	(7.24)	(7.77)
Portfolio parameters	$R_{t-1,t}$	1.52	0.97	0.71	0.55	0.41	0.29	0.17
		(0.82)	(0.52)	(0.41)	(0.36)	(0.35)	(0.37)	(0.42)
	$R^f_{t-1,t}$	1.01	1.85	2.25	2.46	2.70	2.88	3.05
		(0.81)	(0.51)	(0.40)	(0.36)	(0.37)	(0.41)	(0.49)
	$g_{t-1,t}$	-0.10	-0.07	-0.06	-0.05	-0.04	-0.04	-0.03
		(0.03)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)
	$\pi_{t-1,t}$	-0.15	-0.12	-0.11	-0.10	-0.09	-0.09	-0.09
		(0.05)	(0.03)	(0.03)	(0.02)	(0.02)	(0.03)	(0.03)
Optimal portf. as % of Consumption		12.37	8.21	6.68	5.84	5.30	4.91	4.63
Sargan test (*dof = 10*)		12.68	12.70	12.91	12.81	12.36	11.92	11.31
		(0.24)	(0.24)	(0.23)	(0.23)	(0.26)	(0.29)	(0.33)

NOTE: The results in panel (a) are obtained by setting $\rho_{u\eta}=0.5$, those in panel (b) by setting $\rho_{u\eta}=-0.5$. Each column is computed assuming isoelastic preferences for different levels of risk aversion and assuming an annualized discount rate equal to 0.97. The first row of each table reports the estimates of the bound to the fixed annualized participation costs in dollars of year 2000. The next set of rows in the tables contains the estimates of the parameters of the optimal asset portfolio, which implies investing in the financial market the share of time *t* consumption reported in the row before the last. The shares reported are average values. Standard errors in parentheses. The Sargan test of overidentifying restrictions is reported in the last row, with p-values based on 10 degrees of freedom in parentheses.

Table 7: "Cash-in-Hand" Based Estimates of the Lower Bounds to the Transaction Costs for the Market for the Risky Asset and of the Corresponding Optimal Portfolios (17,116 households)

Panel (a): ($\rho_{u\eta}$=0.5)

		γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound		148.26 (30.73)	137.80 (20.04)	140.18 (15.73)	144.73 (13.25)	149.89 (11.90)	155.42 (12.21)	160.21 (15.42)
Portfolio parameters	$R_{t-1,t}$	1.18 (0.57)	0.69 (0.31)	0.48 (0.22)	0.36 (0.19)	0.33 (0.20)	0.22 (0.25)	0.29 (0.32)
	$R^f_{t-1,t}$	0.85 (0.58)	1.45 (0.31)	1.68 (0.23)	1.80 (0.21)	1.81 (0.24)	1.88 (0.33)	1.72 (0.44)
	$g_{t-1,t}$	-0.06 (0.02)	-0.04 (0.01)	-0.02 (0.01)	-0.01 (0.01)	-0.01 (0.01)	0.00 (0.01)	0.01 (0.02)
	$\pi_{t-1,t}$	-0.11 (0.04)	-0.07 (0.02)	-0.05 (0.02)	-0.04 (0.02)	-0.02 (0.02)	-0.01 (0.03)	0.01 (0.04)
Optimal portf. as % of Consumption		16.27	12.53	11.21	10.57	10.20	10.00	9.91
Sargan test (*dof = 10*)		13.18 (0.21)	12.96 (0.23)	12.67 (0.24)	12.71 (0.24)	12.78 (0.24)	11.57 (0.31)	10.40 (0.41)

Panel (b): ($\rho_{u\eta}$=-0.5)

		γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound		148.91 (29.94)	139.40 (19.76)	143.88 (15.65)	146.79 (13.21)	152.23 (11.85)	158.09 (11.92)	163.94 (14.61)
Portfolio parameters	$R_{t-1,t}$	1.18 (0.57)	0.68 (0.31)	0.47 (0.22)	0.35 (0.19)	0.33 (0.19)	0.22 (0.23)	0.27 (0.29)
	$R^f_{t-1,t}$	0.84 (0.58)	1.44 (0.31)	1.68 (0.22)	1.80 (0.20)	1.80 (0.22)	1.88 (0.30)	1.76 (0.40)
	$g_{t-1,t}$	-0.06 (0.02)	-0.04 (0.01)	-0.02 (0.01)	-0.01 (0.01)	-0.01 (0.01)	0.00 (0.01)	0.01 (0.01)
	$\pi_{t-1,t}$	-0.11 (0.04)	-0.07 (0.02)	-0.05 (0.02)	-0.04 (0.01)	-0.02 (0.02)	-0.01 (0.02)	0.01 (0.03)
Optimal portf. as % of Consumption		16.31	12.58	11.26	10.63	10.26	10.06	9.97
Sargan test (*dof = 10*)		13.23 (0.21)	12.99 (0.22)	12.75 (0.24)	12.82 (0.23)	12.77 (0.24)	11.20 (0.34)	9.84 (0.45)

NOTE: The results in panel (a) are obtained by setting $\rho_{u\eta}=0.5$, those in panel (b) by setting $\rho_{u\eta}=-0.5$. Each column is computed assuming isoelastic preferences for different levels of risk aversion and assuming an annualized discount rate equal to 0.97. The first row of each table reports the estimates of the bound to the fixed annualized participation costs in dollars of year 2000. The next set of rows in the tables contains the estimates of the parameters of the optimal asset portfolio, which implies investing in the financial market the share of time *t* consumption reported in the row before the last. The shares reported are average values. Standard errors in parentheses. The Sargan test of overidentifying restrictions is reported in the last row, with p-values based on 10 degrees of freedom in parentheses.

Table 8: Estimates of the Lower Bounds to the Transaction Costs for the Market for the Riskless Asset and of the Corresponding Optimal Portfolios (4,561 households)

Panel (a): ($\rho_{u\eta}$=0.5)

		γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound		24.22	30.99	37.00	42.61	48.55	55.37	63.37
		(3.38)	(4.55)	(5.72)	(7.05)	(9.10)	(12.79)	(19.64)
Portfolio parameters	$R_{t-1,t}$	0.56	0.47	0.49	0.46	0.42	0.38	0.36
		(0.33)	(0.40)	(0.44)	(0.49)	(0.55)	(0.65)	(0.79)
	$R^f_{t-1,t}$	1.75	1.96	1.98	2.03	2.07	2.09	2.05
		(0.42)	(0.49)	(0.54)	(0.58)	(0.65)	(0.74)	(0.88)
	$g_{t-1,t}$	0.01	0.01	0.02	0.03	0.03	0.04	0.05
		(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)	(0.04)
	$\pi_{t-1,t}$	-0.00	0.02	0.03	0.04	0.05	0.07	0.09
		(0.04)	(0.05)	(0.05)	(0.06)	(0.06)	(0.07)	(0.08)
Optimal portf. as % of Consumption		7.79	6.43	5.87	5.52	5.27	5.08	4.92
Sargan test (*dof = 10*)		13.70	14.40	14.46	14.40	14.24	14.01	13.74
		(0.19)	(0.16)	(0.15)	(0.16)	(0.16)	(0.17)	(0.19)

Panel (b): ($\rho_{u\eta}$=-0.5)

		γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound		25.13	32.27	38.88	45.36	52.53	61.23	71.94
		(3.41)	(4.53)	(5.69)	(7.02)	(8.93)	(12.16)	(17.90)
Portfolio parameters	$R_{t-1,t}$	0.29	0.40	0.41	0.35	0.26	0.17	0.28
		(0.30)	(0.37)	(0.41)	(0.44)	(0.49)	(0.56)	(0.67)
	$R^f_{t-1,t}$	2.06	2.02	2.05	2.14	2.23	2.31	2.31
		(0.39)	(0.45)	(0.49)	(0.53)	(0.58)	(0.65)	(0.76)
	$g_{t-1,t}$	0.01	0.02	0.03	0.03	0.04	0.05	0.06
		(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)	(0.04)
	$\pi_{t-1,t}$	-0.01	0.02	0.03	0.03	0.04	0.05	0.07
		(0.04)	(0.04)	(0.05)	(0.05)	(0.05)	(0.06)	(0.06)
Optimal portf. as % of Consumption		7.76	6.43	5.90	5.59	5.38	5.23	5.11
Sargan test (*dof = 10*)		13.22	14.45	14.45	14.33	14.08	13.71	13.37
		(0.21)	(0.15)	(0.15)	(0.16)	(0.17)	(0.19)	(0.20)

NOTE: The results in panel (a) are obtained by setting $\rho_{u\eta}=0.5$, those in panel (b) by setting $\rho_{u\eta}=-0.5$. Each column is computed assuming isoelastic preferences for different levels of risk aversion and assuming an annualized discount rate equal to 0.97. The first row of each table reports the estimates of the bound to the fixed annualized participation costs in dollars of year 2000. The next set of rows in the tables contains the estimates of the parameters of the optimal asset portfolio, which implies investing in the financial market the share of time *t* consumption reported in the row before the last. The shares reported are average values. Standard errors in parentheses. The Sargan test of overidentifying restrictions is reported in the last row, with p-values based on 10 degrees of freedom in parentheses.

Table 9: Estimates of the Lower Bounds to the Transaction Costs for a Portfolio of Risky and Riskless Assets (4,561 households)

Panel (a): ($\rho_{u\eta}$=0.5)

	γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound	60.57 (20.83)	54.86 (14.96)	56.07 (14.61)	58.91 (14.95)	62.90 (16.96)	68.13 (21.19)	74.64 (29.51)
Risky assets (% of Consumption)	12.78	8.73	7.27	6.48	5.96	5.59	5.30
Riskless assets (% of Consumption)	0.00	0.00	0.00	0.00	0.01	0.01	0.01
Sargan test (*dof = 10*)	12.61 (0.25)	12.82 (0.23)	13.12 (0.22)	13.34 (0.21)	12.76 (0.24)	13.93 (0.18)	13.88 (0.18)

Panel (b): ($\rho_{u\eta}$=-0.5)

	γ=0.5	γ=1	γ=1.5	γ=2	γ=2.5	γ=3	γ=3.5
Transaction costs bound	63.42 (22.43)	57.38 (16.10)	58.87 (15.21)	62.21 (16.73)	67.08 (6.95)	73.85 (20.70)	82.76 (28.35)
Risky assets (% of Consumption)	12.72	8.71	7.30	6.53	6.06	5.73	5.48
Riskless assets (% of Consumption)	0.00	0.00	0.00	0.00	0.01	0.01	0.01
Sargan test (*dof = 10*)	12.52 (0.25)	12.78 (0.24)	13.21 (0.21)	13.46 (0.20)	12.76 (0.24)	13.62 (0.19)	13.52 (0.20)

NOTE: The results in panel (a) are obtained by setting $\rho_{u\eta}=0.5$, those in panel (b) by setting $\rho_{u\eta}=-0.5$. Each column is computed assuming isoelastic preferences for different levels of risk aversion and assuming an annualized discount rate equal to 0.97. The first row of each table reports the estimates of the bound to the fixed annualized participation costs in dollars of year 2000. Standard errors in parentheses. The next set of rows in the tables contains the average shares of time *t* consumption to be invested in risky assets and in riskless ones. The estimates of the portfolio parameters are available upon request. The Sargan test of overidentifying restrictions is reported in the last row, with p-values based on 10 degrees of freedom in parentheses.

Figure 1: Expected Utility and Cost Bounds

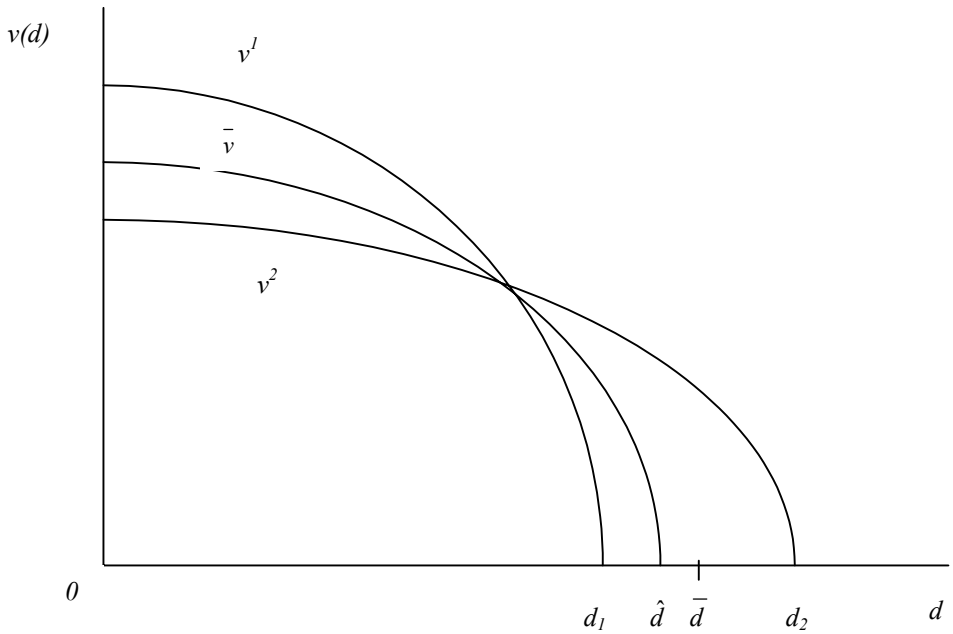

v^h is household h expected utility gain; d_h is the minimal cost which equalises household h expected gain to zero; $\bar{d}$ is the mean household cost bound; $\bar{v}$ is the mean expected gain, which $\hat{d}$ equalizes to zero.