

The Forgone Gains of Incomplete Portfolios

Monica Paiella

Bank of Italy, Research Department, Via Nazionale 91, Rome 00184, Italy

This article proposes a test for the cost-based explanation of nonparticipation, by estimating a lower bound to the forgone gains of incomplete portfolios; these are in turn a lower bound to the costs that could rationalize nonparticipation in financial markets: high bounds would imply implausibly high costs. Assuming isoelastic utility and a relative risk aversion of three or less, for the stock market I estimate an *average* lower bound of between 0.7 and 3.3 percent of consumption. Since total annual (observable plus unobservable) participation costs are likely to exceed these bounds, the cost-based explanation is not rejected by this test. (*JEL* G11, D12 and E21)

Many studies suggest that observed asset returns are inconsistent with consumption choices as predicted by the theory of the intertemporal allocation of consumption. Recent work has explored the possibility that limited participation in financial markets might explain the disparity between theoretical predictions and empirical evidence. More precisely, since the first-order conditions of asset pricing models hold with equality only for those households that own complete portfolios, the models should be tested for this subset and not for the whole population. In practice few households hold shares directly, and therefore the use of aggregate consumption data to evaluate asset pricing models could be very misleading.¹ These points are stressed by Mankiw and Zeldes (1991), Attanasio et al. (2002), Vissing-Jørgensen (2002) and Paiella (2004), among others. While their studies have been somewhat successful in reconciling the consumption-based capital asset pricing model with the empirical evidence, they take limited participation as given and make no attempt to rationalize it.

Limited participation is itself a puzzle for the intertemporal consumption model. Merton (1969) and Samuelson (1969) have illustrated how such behavior is inconsistent with the maximization of expected lifetime utility. Expected-utility maximizers should always be willing to invest

The views expressed are those of the author and do not necessarily reflect those of the Bank of Italy. I am grateful to Orazio Attanasio for his advice and suggestions. I have also benefited from comments by Richard Blundell, John Cochrane, Janice Eberly, Dario Focarelli, John Heaton, Erzo Luttmer, Andrea Tiseno, Annette Vissing-Jørgensen, two anonymous referees, the editor, Maureen O'Hara and participants to the seminars at the University College London, Institute for Fiscal Studies, Bank of Italy, and the NBER Summer Institute on the Aggregate Implications of Microeconomic Behavior. Address correspondence to Monica Paiella, Bank of Italy, Research Department, Via Nazionale 91, Rome 00184, Italy; Tel: +39.06.4792.2595; Fax: +39.06.4792.3723, or e-mail: monica.paiella@bancaditalia.it.

¹ Moreover, since the proportion of households holding shares increased dramatically in the 1980s and 1990s, composition effects might be important in the sub-sample of shareholders.

an arbitrarily small amount in all assets with a positive expected return, including risky ones, unless there are nonlinearities in the budget constraint. Yet, more than two-thirds of the households surveyed in the US Consumer Expenditure Survey for the period 1982–95 did not hold either stocks or bonds.² This behavior is usually reconciled with the intertemporal consumption model by invoking nonproportional costs of financial market participation (explicit and non-explicit). As such costs are for the most part unobservable, the plausibility of the explanation depends crucially on how large they must be to explain the observed data.

This article focuses on the issue of limited participation in financial markets and proposes a heuristic test of the plausibility of the “participation cost explanation.” Using the empirical implications of the consumption model, I estimate a lower bound to the forgone gains of holding an incomplete portfolio. The forgone gains are in turn a lower bound to the level of costs that rationalizes nonparticipation and reconciles incomplete portfolios with the intertemporal consumption model. If the estimated bound is very high, the cost-of-participation explanation will be implausible because costs are unlikely to exceed such level in reality.

My bound estimate varies with risk aversion and is increasing in consumption and wealth. Given a relative risk aversion of three or less, which is the range considered plausible by the literature on the equity premium puzzle,³ and averaging across *all* riskless asset holders who do not hold stocks, I find that the lower bound to the forgone gains for not investing in risky assets ranges from 0.7 to 3.3 percent of expenditure on nondurable goods and services. However, for the wealthiest third of riskless asset holders, the forgone gains could be as high as 6.7 percent. The average bound to the forgone gains for not investing in riskless assets ranges from 0.3 to 1.1 percent.

Overall, the bounds are quite accurately estimated. Nevertheless, since this study uses micro data for consumption, measurement error may be a serious problem for it can make household consumption appear less smooth than it actually is. The less smooth consumption is, the greater are the benefits from trading securities, which can then be used to smooth consumption. Thus measurement error may lead to an overstatement of the benefits of investing in the “overlooked” security and is a potential source of upward bias for my estimates of the forgone gains. A Monte Carlo analysis of its implications suggests that it is indeed a source of

² Similar figures are obtained from the Survey of Consumer Finance.

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

bias. However, the bias appears to be fairly small, especially as regards investments in risky assets and low levels of risk aversion. As expected, it is increasing in risk aversion as more risk averse investors assign greater value to consumption smoothing. More importantly, measurement error biases the bound estimates upwards. Hence, the true bounds can be expected to be lower than the ones I find.⁴

As to the true costs of participation in financial markets, which include both a monetary and nonmonetary component, in 1996, according to the US National Income and Product Accounts, monetary charges for brokerage and trust services and investment counseling amounted to 2.9 percent of consumption (3.7 percent in 2000).⁵ Nonmonetary charges refer to the total value of time spent acquiring and processing information and managing the optimal portfolio. Given an average ratio of consumption of nondurable goods and services to income of 0.5, spending—or just expecting to spend—one hour a week to oversee stock market investments would add up to 5 percent of consumption.⁶ This suggests that the true cost of participating in the stock market is likely to be greater than my estimates of the average bound, supporting the claim that the heuristic test does not reject the “participation cost hypothesis” even at a reasonably high degree of confidence. However, the heterogeneity of the gains and the fact that the wealthiest nonparticipants appear to forgo substantially larger returns imply that caution is needed when using the participation cost hypothesis to explain nonparticipation by individuals with enormous financial wealth.

The estimation of the lower bound to the forgone gains is based on the necessary conditions for the optimality of observed behavior of nonparticipants. The methodology relies on the (empirical) distinction between the consumption paths of households with a well-diversified portfolio and those of individuals with incomplete portfolios. This approach builds on a article on adjustment costs and asset pricing by

⁴ The Monte Carlo simulation is based on the assumption that measurement error is multiplicative and independent across households and from the level of consumption and asset returns, and that it is log-normally distributed with mean one. Its variance is computed as the difference between the time-series variance of average consumption of CEX nonparticipants and the variance of aggregate US expenditure, which is my proxy for measurement-error-free consumption. The variance numbers used in the simulation range from 0.09 to 0.2025. The simulation suggests that, with log-utility, my estimator is biased upward by around 5 percent; the bias goes up to 20 percent when risk aversion is set equal to three and to 31 percent when it is equal to eight.

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

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

Luttmer. Luttmer (1999) focuses on the losses from not exploiting some trading opportunities and proposes a lower bound on the level of fixed transactions costs reconciling per-capita expenditure and asset returns. The forgone gains identified by Luttmer bound from below the cost that a representative agent must pay to trade and to modify her consumption path. I use individual level data instead and, by distinguishing between holders and nonholders of risky assets, I focus on the loss from missing out on the equity premium. Further, by distinguishing between holders and nonholders of riskless assets I also identify the loss from missing out on the riskless rate incurred by those who use other means (such as durables, currency or other) to smooth consumption over time. My frictions are therefore the costs that individual agents must pay in order to participate in financial markets.

The participation costs that I bound from below can be interpreted in two ways. First, they can be thought of as reflecting the information and transactions costs that would induce households not to invest in certain securities. Second, they can be seen as the costs of following “near-rational” decision rules. The idea is that, in practice, households do not literally maximize their utility; they follow heuristic decision processes that economists model as solutions to a maximization problem. Households’ decisions deviate from the solutions to this problem, but with limited costs in terms of utility (for a thorough discussion of this view, see Simon, 1978).

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

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

1. Measuring the Forgone Gains

The framework developed in this section builds on the idea that consumers holding incomplete portfolios could improve their consumption path by investing in assets they do not own. However, this would entail gathering information, making decisions, paying brokerage fees and other costs, all of which acts as a disincentive to complete portfolio diversification. Frictions like these may prevent consumers from exploiting the equity premium paid by risky assets, or even from smoothing consumption by exploiting the riskless rate paid by a riskless asset. The article tests this hypothesis by bounding the costs that would rationalize incomplete portfolios under standard assumptions regarding preferences. The bound is determined by comparing the utility of choosing not to invest in some asset with the utility of the investment, using a revealed preference argument as follows.

Consider an environment in which households have rational expectations, intertemporally, additively separable preferences over consumption, a strictly increasing and concave per-period utility function, $U(\cdot)$, and a positive subjective discount rate, β . Households have access to several means to substitute consumption over time. In particular, they can accumulate financial securities, currency and/or real assets. Some of these substitution opportunities are easier to use than others and we can distinguish three types of households according to portfolio composition: households that hold both risky and riskless financial assets (type 1); those that hold only riskless assets (type 2); and those holding neither (type 3). Let $\{c^h\}_t, t = 1, 2, \dots$ be the observable sequence of consumption choices by household h . These choices are the result of a potentially complicated and unobservable set of decisions involving labor supply, savings and portfolio allocation. However, by assumption, households behave optimally, conditional on the information available. Therefore, the expected gain from choosing any other feasible sequence of consumption bundles, alternative to $\{c^h\}_t$, must be nonpositive.

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

$$\tilde{c}_{h,t} = c_{h,t} - x_{h,t}^c c_{h,t} - \delta^* c_{h,t}, \quad (1)$$

$$\tilde{c}_{h,t+1} = c_{h,t+1} + x_{h,t}^c c_{h,t} R_{t+1} + x_{h,t}^w W_{h,t} (R_{t+1} - R_{t+1}^f), \quad (2)$$

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

$$E_t\{U(\tilde{c}_{h,t}) + \beta U(\tilde{c}_{h,t+1})\} \leq E_t\{U(c_{h,t}) + \beta U(c_{h,t+1})\}. \quad (3)$$

Equation (3) says that, net of the cost δ^* , the expected-utility gain from perturbing the observed consumption path is nonpositive. The investment is thus worthless.

For ease of notation, the forgone utility gain of household h , for the misallocation over the periods t and $t + 1$, is written as:

$$v_{h,t+1}(x_{h,t}^c, x_{h,t}^w, \delta^*) = U(\tilde{c}_{h,t}) + \beta U(\tilde{c}_{h,t+1}) - U(c_{h,t}) - \beta U(c_{h,t+1}). \quad (4)$$

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

$$E_t\{v_{h,t+1}(x_{h,t}^c, x_{h,t}^w, \delta^*)\} \leq 0. \quad (5)$$

If long time series of individual consumption were available, one could proceed as Luttmer (1999) and study (5) by taking unconditional expectations. Instead, the data I use provide only one observation on forgone gains for each nonparticipating household. To deal with this complication, let H_t denote the set of agents, observed at time t , who choose not to invest in the asset considered, based on the information available at t . Then, averaging the individual $v_{h,t+1}(x_{h,t}^c, x_{h,t}^w, \delta^*)$ over H_t , and taking unconditional expectations yields:

$$E \left\{ H_t^{-1} \sum_{h \in H_t} v_{h,t+1}(x_{h,t}^c, x_{h,t}^w, \delta^*) \right\} \leq 0. \quad (6)$$

This inequality holds for any $x_{h,t}^c$ and $x_{h,t}^w$. Therefore, varying $x_{h,t}^c$ and $x_{h,t}^w$ one obtains:

$$\sup_{x_{h,t}^c, x_{h,t}^w} E \left\{ H_t^{-1} \sum_{h \in H_t} v_{h,t+1}(x_{h,t}^c, x_{h,t}^w, \delta^*) \right\} \leq 0. \quad (7)$$

The left-hand-side of (7) is nonnegative at $\delta^* = 0$, and continuous and decreasing in δ^* —as $U(\cdot)$ is continuous and increasing. Hence, we can

construct a lower bound d for δ^* by solving:

$$\sup_{x_{h,t}^c, x_{h,t}^w} E \left\{ H_t^{-1} \sum_{h \in H_t} v_{h,t+1}(x_{h,t}^c, x_{h,t}^w, d) \right\} = 0, \quad (8)$$

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

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

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

2. Estimation Strategy

2.1 Empirical specification and estimation procedure

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

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

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

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

where D_j denotes the derivative with respect to the j^{th} element of $v_{h,t+1}(x_t^c(\cdot), x^w(\cdot), d)$. Equation (9)⁷ determines the optimal investment in case of participation, given the cost. Since, in practice, the actual cost δ^* is neither observed nor estimated, and only a lower bound to the cost is identified, the optimal portfolio is determined as a function of a cost equal to its estimated bound, d , which is consistent with the rest of the analysis. Equation (10) determines the lower bound d to the participation cost δ^* , given the optimal investment.

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

⁷ These are first-order conditions, which are necessary but not sufficient for a maximum, unless the function being maximized is strictly concave in the parameters, which need not be the case in the problem considered here. Thus, second-order conditions must be checked as well.

estimation:

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

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

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

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

$$U(c_{h,t}) = \frac{c_{h,t}^{1-\gamma} - 1}{1-\gamma}, \quad (13)$$

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

In practice, on the basis of the data available, I can identify the bounds to the costs that would rationalize two types of incomplete portfolios. First, I consider the forgone gains for not holding risky assets, but only riskless ones (type 2 households), and bound the cost of not investing in a well-diversified portfolio of risky assets. Second, I look at the forgone gains of type 3 households. Using the information on these households, I can determine a bound to two different costs: the cost that would justify the choice of not investing even in riskless assets alone, and the cost that would justify the choice not to invest in an optimally determined portfolio of risky and riskless assets.⁸

⁸ In principle, when bounding the costs of investing in risky assets I could also consider type 3 households, rather than only type 2. However, the forgone benefits of not investing in risky assets for type 3 households are likely to be very different from those of type 2 households because ultimately the forgone gains also depend on whether the household has other assets. For type 2 households, much of the gains of investing in risky assets can be expected to come from the exploitation of the equity premium. Instead, for type 3 households, they can be expected to be related to consumption smoothing, although at a somewhat high risk. Averaging the two types of gains would make the estimated d uninformative and more difficult to interpret.

2.2 Data

The data are taken from the US Consumer Expenditure Survey (CEX), which covers a representative sample of the US population, and is run on an ongoing basis by the US Bureau of Labor Statistics. The CEX has a rotating panel dimension, with each consumer unit being interviewed every 3 months over a 12-month period. As households complete their participation, new ones are introduced into the panel on a regular basis and, as a whole, about 4500 households are interviewed each quarter, more or less evenly spread over the 3 months.

At the time of the last interview, households provide information on their asset holdings at that date and on the “dollar difference” with respect to the amounts held 12 months earlier. The asset categories in the CEX are: (1) checking, brokerage and other accounts; (2) saving accounts; (3) US savings bonds; (4) stocks, bonds, mutual funds and other securities. As a measure of risky asset holdings, I take the amounts held in “stocks, bonds, mutual funds and other securities” and “US savings bonds”; as a measure of riskless asset holdings, I take the amounts held in checking and saving accounts. In order to determine the asset holding status at t , for each asset category, I subtract from the stocks held at the time of the last interview ($t + 1$) the amount of savings (the dollar change) in the previous 12 months. Table 1 reports the sample composition in each of the years considered on the ground of household asset portfolios. On average, 31 percent of the sample holds positive amounts of both risky and riskless assets; 50 percent holds only riskless assets; 19 percent holds neither. 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 of the survey, which is consistent with the evidence from other datasets.

Each quarterly interview collects detailed information on expenditure for each of the preceding 3 months. However, since the information on asset holdings is annual and I consider household portfolios at the time they enter the survey, for each household I can define only one observation on the expected-utility gain. I therefore use only two observations on consumption, with $c_{h,t}$ and $c_{h,t+1}$ denoting spending in the months preceding the first and the last interview, respectively.⁹ The consumption measure that I use is de-seasonalized, real, monthly, per-adult equivalent

⁹ This timing implies a 9-month gap between t and $t + 1$ and observations on consumption growth over 9 months. For the estimation I therefore use the return on investments of 9 months. The bound to the gains from a 9-month investment are then multiplied by 4/3 to recover the bound to the gains from an investment of 12 months.

Table 1
Sample composition

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

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

expenditure¹⁰ on nondurable goods and services. For consistency, financial wealth is also re-scaled to real, per-adult equivalent terms and is divided by 12 for comparability with the (monthly) consumption measure.

The data used for the analysis cover the period 1982–1995.¹¹ Around 1985–1986, the sample design and the household identification numbers were changed and after the first quarter of 1986 no track was kept of the households that had entered the survey in 1985. As a consequence of this and of the fact that the information on financial asset holdings was collected during the last interview, the households that had their first interview in the third and fourth quarter of 1985 had to be excluded from the sample. The sample used thus consists of households whose first interview was held between 1982:1 and 1985:6 and between 1986:1 and 1995:1 and t runs for a total of 150 periods.

From the initial sample of households, I exclude those with incomplete income responses and those whose financial supplement contains invalid

¹⁰ Nominal consumption is deflated by means of household specific indices based on the Consumer Price Index provided by the Bureau of Labor Statistics. The individual deflators are determined as geometric averages of elementary regional price indices, weighted by the shares of household expenditure on individual goods. See Attanasio and Weber (1995) for a more extensive discussion of these indices. Household per-adult equivalent consumption is obtained from total household consumption using the following adult equivalence scale: the household head is weighted 1, the other adults and the children are weighted 0.8 and 0.4, respectively. By using per-adult equivalent consumption it is possible to account for changes in consumption deriving from changes in household composition occurring between the interviews.

¹¹ Over the period considered stock returns were abnormally high. This implies that any forgone gain based on data for these years tends to overstate expected benefits.

blanks in either the stocks of assets held or in the dollar changes with respect to the previous year. I also exclude those living in rural areas or in university housing, those whose head is younger than twenty-five or older than sixty-five and those who do not participate in all the interviews (about 33 percent of the initial sample). I then select out the top 0.1 percent of the income distribution and the bottom 1.7 percent, which corresponds to about 500 households whose total after-tax annual income is below \$3500 and who are likely to consume all their income, leaving nothing to invest in financial markets. More importantly, these households are likely to be financially constrained and the standard conditions for optimality of behavior, upon which my analysis builds, do not hold for them. I also exclude households with average monthly per-adult equivalent consumption lower than \$250 (about 1000 households corresponding to 3.6 percent of the sample) and those reporting a change in per-adult equivalent consumption, $\Delta c_{h,t}$, greater than \$1750 in absolute value (about 500 households). Finally, I drop the households that hold only risky assets (less than 0.4 percent of the sample). The results are largely unaffected by reasonable changes in the cut-offs. Overall, the sample consists of 23,970 households.

Table 2 reports some descriptive statistics for the sample as a whole and for the three types of households. Type 1 households, who hold both risky and riskless assets, are more likely to be headed by a man, the household head is more educated than the average, slightly older and more often married. Their after-tax income and consumption are also higher. Those holding neither risky nor riskless assets (type 3) tend to be the least educated and to have the lowest income and consumption and in 42 percent of the cases are headed by a woman.

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

	Type 1	Type 2	Type 3	Total sample
Age*	44.1	42.4	43.6	43.1
Less than high school (%)	6.0	15.1	33.7	15.9
High school diploma (%)	50.7	59.0	53.0	55.3
College degree (%)	43.4	25.9	13.4	28.8
Male (%)	77.7	69.4	58.1	70.0
Single person (%)	14.6	20.7	18.3	18.4
Married (%)	76.3	63.3	52.6	65.2
Children (%)	47.5	46.4	52.8	48.0
Living in the Northeast (%)	22.9	19.0	27.3	21.8
Living in the Midwest (%)	26.6	24.6	23.2	24.9
Living in the South (%)	29.4	31.6	33.0	31.2
Living in the West (%)	21.0	24.9	16.7	22.1
After-tax annual household income*	\$70,300	\$47,700	\$37,600	\$51,000
Household financial wealth*	\$49,700	\$11,000	\$0	\$21,300
Annual household consumption*	\$22,700	\$18,100	\$15,400	\$18,900
No. of observations	7,388	12,021	4,561	23,970

Note: Income, financial wealth and consumption are in dollars of year 2000. Consumption consists of spending on nondurables and services. Figures obtained using sample weights. * denotes means.

Table 3
Annual risky and risk-free returns (1981:12–1995:09)

	Mean	Standard deviation	Min	Max
S&P500 composite index	0.122	0.170	−0.277	0.759
Three-month treasury bills	0.026	0.019	−0.007	0.076

Note: Real returns.

Real annual asset returns are summarized in Table 3. The risky return corresponds to the total return (capital gains plus dividends) on the S&P500 Composite Share Index. The riskless return coincides with the return on 3-month US Treasury bills. The mean equity premium that those who do not invest in the risky asset forgo is about 9 percentage points.

3. The Forgone Gains for not Investing in Risky Assets

Tables 4 through 7 present estimates of the gains that riskless asset holders forgo when they choose not to invest in a well-diversified risky portfolio whose return mimics that on the S&P500 CI. The share of consumption to be invested in the asset market after paying the cost, $x_t^c(z_t; a_c)$, is assumed to be a function of a vector of instruments (z_t) that have been shown to be useful in predicting asset returns (see, for example, Keim and Stambaugh (1986) and Fama and French (1989)). The instruments include the returns on the S&P500 CI and on 3-month T-bills, the term spread, and the price-earnings ratio, plus a constant. Variables are lagged one period and refer to the time interval denoted as ‘ $t - 1$ –to– t ’. The tables have the following basic structure. Each column is computed assuming isoelastic preferences for different levels of risk aversion. The rate of discount, β , is set equal to 0.98.¹² Tables 4 and 5 are obtained by averaging the forgone gains over all riskless asset holders. In Tables 6 and 7, households are sorted according to their initial consumption and wealth so that the extent of the differences in the gains related to these characteristics can be appraised.

Table 4 is based on the assumption that riskless asset holders do not move any of their riskless wealth into the risky asset ($x^w(a_w) = 0$). Panel (a) considers the case in which households do not attempt to time the market when determining their optimal investment, which implies a constant trading rule for consumption. When risk aversion is equal to one, the optimal consumption share to invest is around 7 percent of consumption. Consistent with the literature on portfolio choice, this share is decreasing in risk aversion and drops to 4 percent when γ is 3 and to under 3 percent when γ is 10. For a riskless asset holder with a relative risk aversion of one who invests optimally, the forgone gains amount to at least

¹² Higher rates imply slightly higher estimates of d , but the overall conclusions do not change significantly.

Table 4
The forgone gains for not investing in risky assets, with no wealth reallocation

Panel (a)								
RRA	1	2	3	4	5	6	8	10
x^c	0.072 (0.009)	0.050 (0.004)	0.042 (0.003)	0.037 (0.003)	0.034 (0.004)	0.033 (0.006)	0.030 (0.011)	0.027 (0.015)
d	0.006 (0.002)	0.007 (0.001)	0.008 (0.001)	0.010 (0.002)	0.013 (0.003)	0.017 (0.006)	0.027 (0.017)	0.034 (0.029)
Panel (b)								
RRA	1	2	3	4	5	6	8	10
average x_t^c	0.072	0.050	0.041	0.037	0.034	0.032	0.029	0.026
d	0.007 (0.002)	0.008 (0.001)	0.009 (0.001)	0.010 (0.002)	0.013 (0.003)	0.018 (0.006)	0.035 (0.017)	0.054 (0.028)
$x_t^c(z_t; \alpha_c)$								
Constant	-0.601 (1.169)	-0.211 (0.559)	0.258 (0.400)	0.272 (0.377)	0.798 (0.482)	1.209 (0.716)	2.595 (1.143)	3.427 (1.335)
R_t	-4.179 (0.057)	-0.080 (0.028)	-0.056 (0.023)	-0.050 (0.024)	-0.043 (0.028)	-0.043 (0.036)	-0.040 (0.057)	-0.005 (0.102)
R_t^f	0.879 (1.090)	0.358 (0.528)	-0.121 (0.384)	-0.143 (0.360)	-0.642 (0.453)	-1.022 (0.664)	-2.312 (1.052)	-3.115 (1.217)
Term spread $_t$	2.178 (1.194)	1.214 (0.486)	0.934 (0.377)	1.260 (0.460)	1.118 (0.601)	1.252 (0.748)	1.161 (0.905)	0.934 (1.216)
$(P/E)_t$	-0.365 (0.408)	-0.203 (0.172)	-0.274 (0.121)	-0.317 (0.130)	-0.472 (0.184)	-0.638 (0.265)	-1.119 (0.408)	-1.441 (0.439)

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

0.6 percent of expenditure. As γ increases, the point estimate of d rises, although the differences are hardly statistically significant. For values of γ equal to eight or more, the estimate becomes statistically insignificant.

The tendency of the estimated d to increase in γ can be explained on two grounds. First, the benefits of investing in a well-diversified risky portfolio come not just from exploiting the equity premium, but also from the availability of an effective means of smoothing consumption over time, although at a rather high risk. The gains from exploiting the equity premium decrease in risk aversion, because those with a high aversion to risk are *harmed* more by the higher volatility of $c_{h,t+1}$ that an investment in risky assets implies. Instead, the gains from an improved intertemporal reallocation of expenditure increase in risk aversion, because the risk averse value consumption smoothing more. The fact that the estimated d exhibits a tendency to increase in γ suggests that the benefits of smoothing

Table 5
The forgone gains for not investing in risky assets, with wealth reallocation

Panel (a)								
RRA	1	2	3	4	5	6	8	10
x^w	1.00	0.85	0.65	0.50	0.40	0.35	0.25	0.10
x^c	0.049 (0.005)	0.037 (0.002)	0.033 (0.002)	0.031 (0.003)	0.030 (0.004)	0.030 (0.007)	0.029 (0.011)	0.027 (0.015)
d	0.033 (0.010)	0.022 (0.008)	0.018 (0.005)	0.017 (0.004)	0.017 (0.004)	0.020 (0.006)	0.028 (0.017)	0.034 (0.029)
Panel (b)								
RRA	1	2	3	4	5	6	8	10
x^w	1.00	0.85	0.65	0.50	0.40	0.35	0.25	0.25
average x_t^c	0.050	0.037	0.033	0.031	0.030	0.029	0.027	0.025
d	0.033 (0.010)	0.022 (0.008)	0.018 (0.005)	0.017 (0.004)	0.018 (0.004)	0.021 (0.006)	0.036 (0.017)	0.054 (0.028)
$x_t^c(z_t; \alpha_c)$								
Constant	-0.387 (0.638)	0.385 (0.257)	0.322 (0.278)	0.464 (0.361)	0.823 (0.535)	1.456 (0.797)	2.786 (1.268)	3.505 (1.533)
R_t	-0.083 (0.033)	-0.017 (0.016)	-0.019 (0.019)	-0.026 (0.024)	-0.034 (0.031)	-0.037 (0.039)	-0.046 (0.061)	-0.012 (0.110)
R_t^f	0.535 (0.592)	-0.291 (0.237)	-0.236 (0.260)	-0.360 (0.340)	-0.683 (0.501)	-1.267 (0.741)	-2.487 (1.171)	-3.181 (1.397)
Term spread _{t}	1.109 (0.629)	0.580 (0.378)	0.696 (0.426)	0.943 (0.520)	1.161 (0.674)	1.063 (0.832)	1.200 (1.020)	1.012 (1.378)
(P/E) _{t}	-0.211 (0.217)	-0.244 (0.102)	-0.224 (0.116)	-0.303 (0.148)	-0.458 (0.208)	-0.669 (0.296)	-1.165 (0.436)	-1.464 (0.480)

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

consumption with a high-return, well-diversified risky portfolio may be important and large enough to offset the loss due to the higher volatility of $c_{h,t+1}$. Second, some of the increase in the estimated bound is likely to be due to the estimation bias induced by mis-measured consumption. As discussed earlier, measurement error in consumption is a source of upward bias and the bias is increasing in risk aversion.

Panel (b) of Table 4 considers the case in which nonparticipants try to predict excess returns when determining their optimal investment in the risky asset. It turns out that different instruments matter at different levels of risk aversion. In each instance, some are scarcely significant at the standard levels, despite their strong significance in regressions predicting the equity premium. Overall, timing the market appears to increase significantly the gains from participation for the highly risk averse and makes the corresponding bound estimates statistically significant. However, the precision of the estimates for high values of γ is low, as the standard errors are quite high. The bound estimates corresponding to low or intermediate levels of risk aversion are largely unaffected.

The results in Table 5 are based on the assumption that households may find it optimal to move some of their wealth from the risk-free to the risky

asset, once they pay the fixed cost. As before, the overall investment in risky assets is estimated to be decreasing in risk aversion and the forgone gains are statistically nonnegligible even for the most risk averse only if households time the market when determining the optimal trading rule for consumption (panel (b) of the table). The optimal consumption share to invest is decreasing in γ and goes from 5 to 2.5 percent. The optimal fraction of wealth to reallocate turns out to be very high and ranges from 100 percent for the least risk averse to 10 percent.¹³ This result is consistent with the evidence of Heaton and Lucas (1997, 2000), who simulate household portfolio allocation and find that for moderate-to-low levels of risk aversion agents hold only stocks almost all of the time. When households also invest some of their wealth in the stock market, the gains from participation related to the exploitation of the equity premium increase substantially, but so does the volatility of $t + 1$ consumption. Overall, the estimated forgone gain is significantly higher than in the case in which $x^w(a_w) = 0$, but only for households with low risk aversion. When $\gamma = 1$, d is 5 times higher than when $x^w(a_w)$ is set to 0 and equal to 3.3 percent; for $\gamma = 3$, it is twice as large and equal to 1.8 percent, but for $\gamma > 6$, the difference is negligible.¹⁴

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

When $x^w(a_w) = 0$ (Table 6), the bound is increasing in consumption: the higher is consumption, the smaller is the marginal reduction in the time t utility associated with the investment (and the payment of the cost) and this increases the amount households will invest and the overall gain from participation. Generally, the estimated bound is statistically significant only if risk aversion is below five; this could be due to the

¹³ Allowing for households to time the market when choosing the share of wealth to invest in the stock market does not change the results. Tables are available upon request.

¹⁴ Furthermore, if I set $x^w(a_w) = 0.3$, which is the mean portfolio share in stocks in my sample, the estimated gain drops to 1.7 percent when $\gamma = 1$, and to 1.5 when $\gamma = 3$.

Table 6
The forgone gains for different groups of nonstockholders, with no wealth reallocation

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

Note: CRRA utility. RRA is the coefficient of relative risk aversion. *d* denotes the bound estimate as a fraction of time *t* consumption. The trading rule for consumption is specified as: $x_t^c(z_t; a_c) = a_c^c z_t$. Households time the market using past returns on the S&P500 CI and on 3-month Treasury bills, the term spread and the price-earnings ratio. ‘*average x_t^c* ’ denotes the average optimal share of consumption that nonshareholders should invest. The discount rate is set equal to 0.98. (a) The sample consists of 2726 households whose financial wealth is in the lowest third of the sample financial wealth distribution and whose consumption is below the sample median. (b) The sample consists of 1293 households whose financial wealth is in the lowest third of the distribution and whose consumption is above the median. (c) The sample consists of 1957 households whose financial wealth is in the middle third of the distribution and whose consumption is below the median. (d) The sample consists of 2005 households whose financial wealth is in the middle third of the distribution and whose consumption is above the median. (e) The sample consists of 1379 households whose financial wealth is in the top third of the distribution and whose consumption is below the median. (f) The sample consists of 2661 households whose financial wealth is in the top third of the distribution and whose consumption is above the median.

sample sizes, which are fairly small, and to the fact that the standard errors are increasing in γ . When statistically significant, values range between 0.5 and 1.2 percent of consumption.¹⁵ When households can move their wealth from the risk-free to the risky asset, and $x^w(a_w)$ is determined optimally (Table 7), important differences in the size of the forgone gains emerge across the groups considered, with the gains sharply increasing in investors’ resources. Overall, the bound estimate for those in the bottom two-thirds of the financial wealth distribution is under 2 percent. For those in the top third of the distribution, the forgone gains are much greater. For $\gamma = 1$, the bound estimate may be as high as 6.7 percent; for $\gamma = 3$ it drops to around 3.5 percent.¹⁶ These results confirm that higher costs are

¹⁵ An exception are those with large wealth, low consumption and very high risk aversion, whose gains turn out to be much higher.

¹⁶ The gains of those with large wealth and low initial consumption are slightly smaller if they are prevented from short-selling the risky asset, which allows them to smooth consumption over the two periods.

Table 7
The forgone gains for different groups of nonstockholders, with wealth reallocation

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

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

needed to explain the choices of the wealthiest nonparticipants, which is fully consistent with the evidence of Vissing-Jørgensen (2003).

4. The Forgone Gains for not Investing in Financial Assets

Tables 8 and 9 are based on a set of households that hold neither risky nor riskless assets and on the assumption that they smooth consumption by other means, on which no information is available, however. Table 8 focuses on the forgone gains for not investing in an asset yielding the 3-month T-bill rate of return. In panels (a) and (b) households can only *buy* the asset; in panel (c) short-selling, i.e. borrowing at the riskless rate, is allowed. In panel (a) households do not time the market; in panels (b) and (c), they do so using lagged returns on the riskless and risky assets. The results are very similar across the three panels and are robust to using different sets of instruments to predict returns. According to the evidence in the table, a household with no financial assets and a risk aversion of one could increase its utility by investing around 5 percent of its consumption for one period in the riskless asset. As before, as risk aversion increases, the utility maximizing investment decreases. For high values of γ , there

is some evidence that optimal behavior may involve some short-selling of the riskless asset,¹⁷ which may reflect some type of liquidity constraint. At high values of risk aversion, results should therefore be appraised with care. Overall, the estimated forgone gains turn out to be quite small and, as expected, are smaller than those recorded for nonparticipation in the risky asset market. The gains are increasing in risk aversion, but for $\gamma > 5$ they are statistically negligible. For low-to-moderate levels of risk aversion, they range from 0.3 to 2.6 percent of expenditure on nondurable goods and services. Note that allowing for borrowing does not increase the gains in any significant way.

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

$$\tilde{c}_{h,t} = c_{h,t} - x^c(a_c)c_{h,t} - \delta_{r,rf}^*c_{h,t}; \quad (14)$$

$$\tilde{c}_{h,t+1} = c_{h,t+1} + x^c(a_c)c_{h,t}R_{t+1} + y(a_y)c_{h,t}(R_{t+1} - R_{t+1}^f), \quad (15)$$

where $\delta_{r,rf}^*$ denotes the cost that household h has to pay to participate in a market where it can trade risky and riskless assets. $x^c(a_c)$ denotes the fraction of time t expenditure the household is willing to give up and invest in financial assets after paying the cost. $y(a_y)$ determines the allocation between risky and riskless assets. The investments in risky and riskless assets are equal to $(x^c(a_c) + y(a_y))c_{h,t}$ and $-y(a_y)c_{h,t}$, respectively. When $y(a_y) < 0$, both assets are bought in positive amounts.¹⁸ When $y(a_y) = 0$, there is no investment in the riskless asset; when $y(a_y) > 0$, the household borrows at the risk-free rate to invest in the risky asset. The results in panel (a) of Table 9 have been obtained by imposing no borrowing; in panel (b) and (c) borrowing is allowed. In all cases, $x^c(a_c) \geq 0$, which implies that households consume all the return on the investment at time $t + 1$. Allowing the investment functions to depend on the vector of instruments does not change the results in any significant way.

The results in panel (a) of Table 9 suggest that if type 3 households were to participate in financial markets and could choose to invest in either risky or riskless assets or in both, for low-to-moderate levels of risk aversion, their utility maximizing portfolio would consist basically just of risky assets. As before, the consumption share that households find optimal to invest is decreasing in risk aversion and ranges from 7.5 to 2.4 percent. Note that the overall investment would be higher than that in

¹⁷ For $\gamma \geq 4$, the average x_t^c in panel (b) (where short-selling is not allowed) is slightly higher than the average x_t^c in panel (c) (where short-selling is allowed). This is due to the fact that when short-selling is allowed, x_t^c is nonpositive for some t .

¹⁸ $y_t(z_t; a_y) < -x^c(a_c)$ is never optimal as it would imply borrowing at the risky rate to invest in the riskless asset.

Table 8
The forgone gains for not investing in riskless assets

Panel (a) RRA		1	2	3	4	5	6	8	10
x^c	x^c	0.052	0.048	0.045	0.042	0.040	0.036	0.029	0.024
	d_{rf}	0.003	0.007	0.011	0.016	0.025	0.035	0.048	0.048
		(0.001)	(0.001)	(0.002)	(0.005)	(0.011)	(0.022)	(0.045)	(0.052)
	$x^c(\alpha_c)$								
	Constant	2.907	2.993	3.062	3.119	3.184	3.273	3.495	3.702
		(0.235)	(0.128)	(0.112)	(0.145)	(0.231)	(0.351)	(0.570)	(0.685)
Panel (b) RRA		1	2	3	4	5	6	8	10
x_t^c	average x_t^c	0.052	0.048	0.045	0.042	0.038	0.033	0.028	0.032
	d_{rf}	0.003	0.007	0.011	0.016	0.026	0.038	0.058	0.063
		(0.001)	(0.001)	(0.002)	(0.005)	(0.012)	(0.024)	(0.048)	(0.056)
	$x_t^c(z_t; \alpha_c)$								
	Constant	1.509	1.639	1.880	1.857	-11.677	-26.752	-116.934	-169.754
		(15.091)	(7.607)	(6.772)	(9.070)	(13.218)	(33.190)	(96.316)	(112.778)
R_t	R_t	0.090	0.109	-0.111	-0.386	-0.684	-0.421	2.373	3.663
		(1.675)	(0.884)	(0.758)	(0.855)	(1.399)	(2.356)	(3.165)	(3.528)
	R_t^f	1.274	1.207	1.272	1.645	15.312	29.998	116.528	167.595
		(15.155)	(7.342)	(6.518)	(8.891)	(12.986)	(31.721)	(93.699)	(109.852)
Panel (c) RRA		1	2	3	4	5	6	8	10
x_t^c	average x_t^c	0.052	0.048	0.045	0.041	0.037	0.031	0.015	0.004
	d_{rf}	0.003	0.007	0.011	0.017	0.026	0.038	0.058	0.066
		(0.001)	(0.001)	(0.002)	(0.005)	(0.012)	(0.024)	(0.046)	(0.050)
	$x_t^c(z_t; \alpha_c)$								
	Constant	0.140	0.229	0.314	0.471	0.675	0.896	1.266	1.476
		(0.813)	(0.389)	(0.306)	(0.352)	(0.458)	(0.555)	(0.564)	(0.422)
R_t	R_t	-0.002	0.002	0.018	0.039	0.065	0.096	0.177	0.252
		(0.088)	(0.045)	(0.039)	(0.049)	(0.069)	(0.098)	(0.148)	(0.136)
	R_t^f	-0.084	-0.180	-0.282	-0.461	-0.692	-0.947	-1.411	-1.706
		(0.832)	(0.389)	(0.307)	(0.364)	(0.485)	(0.601)	(0.647)	(0.421)

Note: CRRA utility. RRA is the coefficient of relative risk aversion. d_{rf} denotes the bound estimate as a fraction of time t consumption. ‘average x_t^c ’ and x^c denote the (average) optimal share of time t consumption to invest in the riskless asset when households do and do not time the market when determining their investment, respectively. The estimates in panel (a) are based on the assumption that households do not time the market when determining their investment and cannot short-sell the riskless asset. The (constant) trading rule for consumption is specified as: $x^c(a_c) = (1 + \exp(a_c))^{-1}$. The estimates in panel (b) are based on the assumption that households do time the market using the past returns on the S&P500 CI (R_t) and on 3-month Treasury bills (R_t^f). The trading rule for consumption is specified as $x_t^c(z_t; a_c) = (1 + \exp(a_c'z_t))^{-1}$. The estimates in panel (c) are based on the assumption that households time the market and are allowed to short-sell the riskless asset. The trading rule for consumption is specified as: $x_t^c(z_t; a_c) = a_c'z_t$. The discount rate is set equal to 0.98. The sample includes 4561 households who own neither risky nor riskless assets.

riskless assets, if the latter were the only asset available (Table 8), which suggests that the substitution effect prevails over the income effect. The estimated forgone gains are increasing in γ and are also generally higher than those for not holding just riskless assets, reported in Table 8, owning

Table 9
The forgone gains for not investing in either risky or riskless assets

Panel (a)								
RRA	1	2	3	4	5	6	8	10
$(x^c + y)$	0.075	0.056	0.049	0.043	0.039	0.035	0.027	0.000
$(-y)$	0.000	0.001	0.001	0.001	0.001	0.001	0.002	0.024
$d_{r,rf}$	0.007	0.010	0.013	0.019	0.027	0.036	0.049	0.048
	(0.001)	(0.001)	(0.002)	(0.005)	(0.011)	(0.022)	(0.045)	(0.052)
Panel (b)								
RRA	1	2	3	4	5	6	8	10
$(x^c + y)$	3.019	1.303	0.902	0.633	0.451	0.312	0.097	0.000
$(-y)$	-3.010	-1.279	-0.877	-0.603	-0.417	-0.278	-0.068	0.024
$d_{r,rf}$	0.127	0.057	0.042	0.036	0.037	0.041	0.049	0.048
	(0.061)	(0.026)	(0.018)	(0.013)	(0.013)	(0.022)	(0.045)	(0.052)
Panel (c)								
RRA	1	2	3	4	5	6	8	10
$(x^c + y_1)$	1.153	0.534	0.318	0.193	0.092			
$(-y_1)$	-1.087	-0.483	-0.271	-0.149	-0.052	No borrowing		
$d_{r,rf}$	0.018	0.014	0.016	0.020	0.027			
	(0.022)	(0.010)	(0.006)	(0.005)	(0.011)			

Note: CRRA utility. RRA denotes the coefficient of relative risk aversion. $d_{r,rf}$ denotes the bound estimate as a fraction of time t consumption. $(x^c + y)$ and $(x^c + y_1)$ denote the optimal share of time t consumption to invest in the risky asset. $(-y)$ denotes the optimal share of time t consumption to invest in the riskless asset. When $(-y) \leq 0$ borrowing occurs. $(-y_1) \leq 0$ indicates borrowing at a rate equal to twice the riskless rate. The trading rule for consumption is specified as: $x^c(a_c) = (1 + \exp(a_c))^{-1}$. The estimates in panel (a) are based on the assumption that households cannot borrow at the riskless rate to finance the risky asset investment and the investment in the riskless asset is specified as $(-y(a_y)) = (1 + \exp(a_y))^{-1}$. The estimates in panel (b) are based on the assumption that they can borrow as much as they want at the riskless rate and $y(a_y) = a_y$. The estimates in panel (c) are based on the assumption that households can borrow at a rate equal to twice the riskless rate and $y_1(a_{y1}) = a_{y1}$. The tables with the estimates of the parameters α_c , α_y and α_{y1} and their standard errors are available upon request.

to the higher return on the overlooked investment. Overall, they range from 0.7 to 2.7 percent and are statistically insignificant for $\gamma > 5$.

The forgone gain estimates turn out to be much larger when households are allowed to take short positions in the riskless asset to invest in stocks, as panel (b) of Table 9 shows. A household with log utility would maximize its participation gains if it could invest almost three times its consumption in the risky asset and finance that investment at the risk-free rate. In this instance, its forgone gains would amount to 13 percent of its consumption. However, as risk aversion increases, the optimal short position decreases rapidly, together with the size of the forgone gains. For $\gamma = 3$, the optimal investment in the risky asset would imply borrowing a sum equal to 88 percent of consumption and the corresponding gain would be 4.2 percent. For $\gamma > 8$, the gains are statistically negligible. Furthermore, it seems unlikely that households would be able to borrow at the 3-month T-bill rate to finance a stock investment. Over the period considered, the mean rate charged on credit card balances, which can be taken as an indicator

of the rates charged on noncollateralized loans, has been 2.9 times the rate on 3-month T-bills.¹⁹ Panel (c) considers the case in which households can finance their risky asset investment by borrowing at a rate set equal to twice the risk-free. In this case, for $\gamma = 1$, the optimal investment in the risky asset would imply borrowing a sum approximately equal to household expenditure and the corresponding gain would be 1.8 percent. When $\gamma \geq 5$, households would not find it optimal to borrow to invest in the risky asset.

5. Concluding Remarks

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

Using the information on households that hold neither risky nor riskless assets, it is possible to have an idea of the forgone gains for not investing in riskless assets and in an optimally determined portfolio of risky and riskless assets. The forgone gain for not investing in riskless assets is increasing in risk aversion and ranges from 0.3 of their consumption, if risk aversion is one, to 2.6 percent, if risk aversion is five (for higher levels of γ , it is statistically insignificant). When allowed to invest in both risky and riskless assets, households have a strong preference for the former. If they were allowed to borrow at the risk-free rate and invest in the risky asset and risk aversion were sufficiently low, they would borrow up to three times their consumption and their forgone gains could be as high as 13 percent. However, such levels of borrowing are unlikely to occur in practice, because it seems unlikely households would be able to borrow at the risk-free rate. With a borrowing rate twice as high as that on 3-month T-bills, the forgone gains would drop to less than 2 percent.

These estimates are a structural implication of the intertemporal consumption model. They provide a natural test of the plausibility of

¹⁹ See the Federal Reserve Board statistical release “G19, Consumer Credit”.

the fixed costs hypothesis as an explanation for limited participation by setting a lower bound on the level of costs needed to rationalize incomplete portfolios. The plausibility of the “cost-of-participation hypothesis” is an empirical issue and depends crucially on the size of the actual costs of financial market participation relative to the theoretical predictions: for the cost-of-participation hypothesis not to be rejected, the actual costs should be higher than my estimates.

Overall, there is limited empirical evidence on the nature and size of the costs associated with financial transactions, especially at the household level. A reason is that some of these costs are likely to be nonobservable and related to information gathering and processing, especially where investing in the stock market is concerned. Given a ratio of consumption of nondurable goods and services to income of 0.5, spending just one hour a week overseeing personal investments would add up to a nonmonetary cost of 5 percent of household expenditure. To this must be added monetary charges, which, based on the US National Income and Product Accounts, could be as much as 3 percent of household consumption.

These figures suggest that the true costs of participating in the stock market are likely to be larger than my estimated bounds, supporting the claim that the heuristic test does not reject the “participation cost hypothesis”, even at a reasonably high degree of confidence. In fact, due to measurement error in consumption data, the “true” bounds are likely to be smaller than the ones I find. Therefore, on the basis of my evidence, the hypothesis that participation costs may reconcile the stock market participation anomaly with the intertemporal choice model cannot be rejected. Nevertheless, the heterogeneity in the gains from participation and the fact that the wealthiest nonparticipants appear to forgo substantially larger returns suggest that caution is needed when considering the participation cost hypothesis as an explanation of nonparticipation by those with substantially large amounts of financial wealth; this is in line with Vissing-Jørgensen (2003) conclusions. Finally, it cannot be ruled out that participation costs may reconcile the choices of households that hold neither risky nor riskless assets. In fact, fairly modest costs are sufficient to rationalize their behavior, given the presence of some constraint on the possibility of financing the investment in risky assets by borrowing at the risk-free rate.

References

- Attanasio, O., J. Banks, and S. Tanner. 2002. Asset Holding and Consumption Volatility. *Journal of Political Economy* 110:771–92.
- Attanasio, O., and G. Weber. 1995. Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy* 103:1121–57.
- Cochrane, J. 1989. The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives. *American Economic Review* 79:319–37.

- Fama, E., and K. French. 1989. Business Conditions and Expected Returns on Stocks and Bonds. *Journal of Financial Econometrics* 25:23–49.
- Heaton, J., and D. Lucas. 1997. Market Frictions, Saving Behavior and Portfolio Choice. *Macroeconomic Dynamics* 1:76–101.
- Heaton, J., and D. Lucas. 2000. Portfolio Choice in the Presence of Background Risk. *Economic Journal* 110:1–26.
- Kandel, S., and R. Stambaugh. 1991. Expectations and Volatility of Consumption and Asset Returns. *Review of Financial Studies* 3:207–32.
- Keim, D., and R. Stambaugh. 1986. Predicting Returns in the Stock and Bond Markets. *Journal of Financial Economics* 17:357–90.
- Kocherlakota, N. 1990. On Tests of the Representative Agent Asset Pricing Models. *Journal of Monetary Economics* 6:285–304.
- Lucas, D. 1994. Asset Pricing with Undiversifiable Risk and Short Sales Constraints: Deepening the Equity Premium Puzzle. *Journal of Monetary Economics* 34:325–42.
- Luttmer, E. 1999. What Level of Fixed Costs Can Reconcile Asset Returns and Consumption Choices? *Journal of Political Economy* 107:969–97.
- Mankiw, N., and S. Zeldes. 1991. The Consumption of Stockholders and Non-Stockholders. *Journal of Financial Economics* 17:211–9.
- Mehra, R., and E. Prescott. 1985. The Equity Premium Puzzle. *Journal of Monetary Economics* 15:145–61.
- Merton, R. 1969. Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. *Review of Economics and Statistics* 51:247–57.
- Newey, W., and K. West. 1987. A Simple Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55:703–8.
- Paiella, M. 2004. Revisiting the Implications of Heterogeneity in Financial Market Participation for the C-CAPM. *Review of Finance* 8:445–80.
- Samuelson, P. 1969. Lifetime Portfolio Selection by Dynamic Stochastic Programming. *Review of Economics and Statistics* 51:39–46.
- Simon, H. 1978. Rationality as Process and as Product of Thought. *American Economic Review* 68:1–16.
- Vissing-Jørgensen, A. 2002. Limited Asset Market Participation and the Elasticity of Intertemporal Substitution. *Journal of Political Economy* 110:825–53.
- Vissing-Jørgensen, A. 2003. Perspectives on Behavioral Finance: Does ‘Irrationality’ Disappear with Wealth? Evidence from Expectations and Actions. NBER Macroeconomics Annual.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.