Review of Finance **8:** 445–480, 2004.
© 2004 *Kluwer Academic Publishers. Printed in the Netherlands.*

Heterogeneity in Financial Market Participation: Appraising its Implications for the C-CAPM ⋆

MONICA PAIELLA
Bank of Italy, Economic Research Department

Abstract. This paper presents empirical evidence that accounting for heterogeneity in financial market participation is important for evaluating the empirical performance of the Consumption-based Capital Asset Pricing Model (C-CAPM). Using the U.S. Consumer Expenditure Survey as a common testing ground, I re-assess three well-known characterizations of the "equity premium puzzle": (i) the inconsistency of the representative agent's IMRS with Hansen and Jagannathan bounds; (ii) Mehra and Prescott's calibration of a large representative agent's risk aversion; (iii) Hansen and the Singleton's large structural estimates of the preference parameters based on aggregate data. In all three cases, the estimates of risk aversion conditional upon financial market participation are not as far from reasonable values as the corresponding unconditional ones. The differences suggest that part of the "equity premium puzzle" can be accounted for by the use of a "representative agent" assumption rather than a more appropriate "representative stockholding agent" assumption.

1. Introduction

The Consumption-based Capital Asset Pricing Model (C-CAPM) by Merton (1973), Breeden (1979), Lucas (1978) and Brock (1982) is a popular framework for understanding the valuation of assets and the serial correlation properties of asset returns and consumption. Its development represents an extension of the Capital Asset Pricing Model (CAPM) of Sharpe (1964) and Lintner (1965) to an intertemporal setting. The C-CAPM is based on the standard representative agent model for consumption, and prices assets by discounting returns using the intertemporal marginal rate of substitution. Specifically, the price of a claim is determined by the product of its expected payoff and the marginal rate of substitution of the representative investor. Empirically, the C-CAPM has not performed well. There are at least three characterizations of this failure. First of all, Mehra and Prescott (1985)'s calibration of the equity premium puzzle has illustrated the impossibility of finding a plausible pair of values for the subjective discount rate and the coefficient of relative risk aversion consistent with the observed risk premium. Secondly, Hansen

⋆ The author would like to thank Orazio Attanasio, James Banks, Luigi Guiso, Marco Pagano, Andrea Tiseno, Giovanni Verga and two anonymous referees. The views expressed are those of the author and do not necessarily reflect those of the Bank of Italy. All errors and omissions are the author's responsibility.

and Jagannathan (1991) determine a bound on the first two moments of a generic stochastic discount factor to assess whether particular parameterizations of asset pricing models are consistent with the observed asset market data. They show that for the representative agent's intertemporal marginal rate of substitution (IMRS) to be consistent with the bound – and therefore with observed returns – the representative investor risk aversion must be very high, around 30 when preferences are isoelastic. Thirdly, Hansen and Singleton (1982, 1984) estimate the structural parameters of the model and obtain implausible values for the preference parameters; furthermore, they statistically reject the set of overidentifying restrictions implied by the C-CAPM. The empirical failure of the C-CAPM is often quoted as evidence against the hypotheses of rational expectations and market efficiency, which lie behind most modern economic and financial models.

There is now a vast literature aimed at rationalizing this evidence. Epstein and Zin (1989, 1991), Constantinides (1990), Heaton (1995) and Campbell and Cochrane (1995), among others, focus on individual preferences, relax the assumption of time-separable power utility, and consider the implications of generalized expected utility, habit formation and relative consumption effects. Another possibility that has been examined, for example by Heaton and Lucas (1996), is to allow for incomplete markets and trading costs. The strengths and weaknesses of these approaches are reviewed extensively by Kocherlakota (1996).

More recently, along the lines of a paper by Mankiw and Zeldes (1991), there has been growing interest in the implications of limited participation in financial markets, which would make per-capita consumption growth a poor proxy for individual consumption growth when pricing assets on the basis of their covariance with the individual intertemporal marginal rate of substitution. The idea that accounting for heterogeneity in participation might explain the disparity between predictions and evidence depends on the presumption that there is no theoretical reason for non-shareholders to adjust their consumption growth in response to predictable changes in share returns (unless their income is correlated with share returns). If non-shareholders' consumption co-varies with returns much less than shareholders', then to include the former in the consumption measure for the C-CAPM will lead to inconsistent estimates and to the rejection of the model's predictions. Further, the fact that a fraction of households hold the entire stock market portfolio might make their consumption and, as a consequence, their stochastic discount factor more volatile than that of non-shareholders.

The objective of this paper is to examine the quantitative implications of limited financial market participation for the three leading characterizations of the empirical failure of the C-CAPM within a common testing ground, based on the U.S. Consumer Expenditure Survey (CEX). I start with Hansen and Jagannathan's framework and compare the IMRSs of shareholders and those of the CEX population-wide representative agent with the Hansen and Jagannathan bound. To appraise the evidence, I take an evaluation criterion that combines the Euclidean distance of the IMRS from the bound and the coefficient of risk aversion. For values of the coefficient of relative risk aversion below 6, such criterion achieves a minimum when pricing assets using the IMRS of the representative shareholder, as opposed to that of the population-wide representative agent. For shareholders, the Euclidean distance of the IMRS from the bound is minimized for a risk aversion around 4. For the population-wide representative agent the minimum is achieved for a much higher value. Next, I calibrate Mehra and Prescott equity premium and estimate the coefficient of relative risk aversion for shareholders and for the population-wide representative agent, exploiting the unconditional restrictions implied by the model. The two coefficients are significantly different and the estimate of risk aversion based on shareholder consumption is about 40 percent lower than that based on the consumption of all households. Last, I take a structural approach and use the conditional distribution of the changes in expectations of consumption and asset returns to estimate efficiently the parameters of interest in the spirit of Hansen and Singleton (1982, 1984). Generalized method of moments estimation of the Euler equation using shareholder data yields sensible and precise estimates of the preference parameters, whereas the risk aversion estimates based on data that include non-shareholder expenditure are significantly larger and statistically imprecise.

Using three different approaches on a single data set to assess the implications of limited participation for the empirical performance of the C-CAPM distinguishes this work from the previous research on this issue, which typically provides evidence based on just one estimation method. The three approaches yield somewhat different point estimates of the preference parameters conditional on participation. All three estimates differ from the corresponding unconditional estimate in the same direction, pointing out that part of the equity premium anomaly illustrated by Mehra and Prescott (1995), Hansen and Jagannathan (1991) and Hansen and Singleton (1982, 1984) can be accounted for by the use of a "representative agent" assumption rather than a more appropriate "representative stockholding agent" assumption. Since the data set is the same throughout the analysis, the differences in the point estimates derive from the fact that the different approaches place a different degree of structure on the problem and rely on more and more stringent hypotheses. The more stringent the hypotheses, the more efficient the estimates, but the less robust the inference. Interestingly and reassuringly, the differences across the estimation methods are relatively smaller when using shareholders' expenditure rather than aggregate expenditure. Further, each of the three approaches emphasizes different features of the data and highlights their role in accounting for the evidence.

The main related studies are Mankiw and Zeldes (1991), Attanasio et al. (2002) and Vissing-Jørgensen (2002). Mankiw and Zeldes (1991) re-calibrate Mehra and Prescott equity premium using the U.S. Panel Study of Income Dynamics (PSID) and find that distinguishing between shareholders and non-shareholders helps explain the size of the equity premium. However, they do not fully resolve the equity premium puzzle. Their estimate of shareholders' risk aversion is substantially lower than that of non-shareholders' (though no standard errors are provided), but it is still too high with respect to the values that are deemed plausible in the literature. Studies by Arrow (1971), Friend and Blume (1975), and Kydland and Prescott (1982), among many others, provide plenty of evidence for restricting the value of relative risk aversion to be below 10, whereas Mankiw and Zeldes' estimate of risk aversion for shareholders is 35 (262 for non-shareholders).

Attanasio et al. (2002) and Vissing-Jørgensen (2002) provide additional support for limited asset market participation theory using the U.K. Family Expenditure Survey and the U.S. Consumer Expenditure Survey, respectively. They employ a conditional version of the model, i.e., exploit the information about predictable movements in expected consumption growth rates and asset returns in the spirit of Hansen and Singleton (1982, 1984). Their estimates of the preference parameters, based on shareholder data only, are sensible and efficient. An important innovation by Attanasio et al. (2002) is the estimation of ownership probabilities to separate "likely" shareholders from non-shareholders, which allows to control for sample selection into the group. This approach, which is used also in the present paper, avoids the potential bias that might be induced by unobserved time-varying taste shocks affecting both consumption and participation choice. Its drawback is that it brings into the analysis some noise related to the misclassification of households. However, the inability to perfectly identify shareholders biases against finding differences in risk aversion estimates for shareholders and the representative agent, as the Monte Carlo simulation reported in the Appendix shows. This implies that the implications of limited financial market participation for the empirical appraisal of the C-CAPM and the benefits of accounting for it are even stronger than those implied by the evidence presented here and in Attanasio et al..

The rest of the paper is organized as follows. In Section 2, I review briefly the C-CAPM and derive its testable implications. Section 3 sets forth the conditions for aggregation that justify the use of aggregate data in the empirical implementation of the model. I then discuss the implications of limited financial market participation for the analysis and the econometric techniques to account for it, when the conditions for aggregation are violated and households are heterogeneous (at the margin). Section 4 presents the data and some evidence on household portfolios and share ownership based on the U.S. Consumer Expenditure Survey. Section 5 assesses the empirical performance of the model when restricted to shareholding households versus its performance under the assumption of population-wide representative agent. For most of the analysis I maintain the (standard) assumption of isoelastic utility; within Hansen and Jagannathan's and Mehra and Prescott's frameworks, I check also the implications of Epstein and Zin-type utility, which is a generalization of the standard isoelastic specification. Section 6 appraises the evidence and concludes.

2. The Consumption-based Capital Asset Pricing Model

Consider the standard intertemporal optimization problem facing a consumer with access to N different assets. The consumer chooses consumption and portfolio by maximizing expected lifetime utility (appropriately discounted) subject to an intertemporal budget constraint that reflects intertemporal allocation possibilities. Assuming that preferences over future random consumption streams are intertemporally additive, the maximization problem for individual h can be written as:

$$\begin{aligned}&\max E_t\left[\sum_{s=t}^{T}\beta^{s-t}U(c_{h,s})\right];\\&\text{s.t.}\\&\sum_{k=1}^{N}A_{h,s}^{k}=\sum_{k=1}^{N}A_{h,s-1}^{k}(1+r_s^k)+y_{h,s-1}-c_{h,s-1},\end{aligned} \tag{1}$$

where $c_{h,s}$ denotes individual h consumption in period s; $y_{h,t}$ is the consumer's non-asset income, β the time-discount factor, $A_{h,t}^k$ the amount of wealth held in asset k, whose return is r_{t+1}^k. If asset k is held at t and $t+1$, the first-order condition for the problem is:

$$U'(c_{h,t})=E_t[\beta U'(c_{h,t+1})(1+r_{t+1}^k)] \tag{2}$$

which can be re-written as:

$$E_t[m_{h,t+1}(1+r_{t+1}^k)]=1, \tag{3}$$

where $m_{h,t+1}=\beta U'(c_{h,t+1})/U'(c_{h,t})$. $m_{h,t+1}$ denotes the marginal rate of substitution (IMRS) between consumption at time t and at $t+1$ and is known as the stochastic discount factor. Since (3) must hold for all assets held in non-zero amounts at t and at $t+1$, it can also be expressed in terms of excess return between, for example, a risky asset k and a risk-free asset f:

$$E_t[m_{h,t+1}(r_{t+1}^k-r_{t+1}^f)]=0. \tag{4}$$

(3) and (4) impose unambiguous statistical restrictions on the co-movement between consumption changes and asset returns.

A key implication of Equation (3) is that equilibrium returns are determined by the intertemporal marginal rate of substitution of consumption. When markets are complete, observed asset returns are sufficient to identify the IMRS. When markets are incomplete, there may be idiosyncratic variation in individual investors' marginal utilities, hence multiple stochastic discount factors that are consistent with (3). However, as Hansen and Jagannathan (1991) show, observed returns can still be used to derive a set of restrictions on the pricing factors. Using only financial data, they restrict the admissible region for the mean and standard deviation of $m_{h,t+1}$ by constructing minimum variance stochastic variables that vary with returns in the same way as the IMRS in the Euler equation in (3).[1] The region provides a convenient summary of the sense in which asset market data are anomalous from the vantage point of intertemporal asset pricing theory.

Alternative tests of (3) and (4) can be designed using a more structural approach and the orthogonality conditions that they imply to estimate the preference parameters. A simple test based on the unconditional restrictions that the model (1) implies can be obtained using the law of iterated expectations and replacing the conditional expectations (3) and (4) with unconditional ones, so that:

$$E[m_{h,t+1}(1 + r^k_{t+1})] = 1, \tag{5}$$

$$E[m_{h,t+1}(r^k_{t+1} - r^f_{t+1})] = 0. \tag{6}$$

Then, given a sufficiently long time series of data on $c_{h,t}$ and asset returns, we can calibrate the expectation on the left-hand side of (6) using the sample mean of:

$$e_{h,t+1} = m_{h,t+1}(r^k_{t+1} - r^f_{t+1}), \tag{7}$$

and determine the preference parameters in $m_{h,t+1}$ that justify observed behavior. The logic behind this test is that only if $E(e_{h,t+1}) = 0$ will the agent be at an optimum and unable to improve his welfare by borrowing at the risk-free rate and investing in the risky asset. Note that Hansen and Jagannathan's lower bound on the volatility of $m_{h,t+1}$ is an even weaker implication of (4) than (6), in that (6) implies the variance bound while the converse is not true.

The relationship upon which Mankiw and Zeldes (1991) rely can be derived from (5) by adding the assumptions of isoelastic utility and of log-normal consumption growth and asset returns. Under these hypotheses, the left-hand-side of (5) can be written as:

$$E[m_{h,t+1}(1 + r^k_{t+1})] = \exp\left\{-\gamma E\left[\log \frac{c_{h,t+1}}{c_{h,t}}\right] + E[\log R^k_{t+1}] + \delta^k_h\right\}, \tag{8}$$

where γ is the coefficient of relative risk aversion and the inverse of the elasticity of intertemporal substitution, R^k_{t+1} is the gross return on asset k and:

[1] In practice, Hansen and Jagannathan (1991) focus on unconditional moment restrictions, because they are typically easier to estimate than conditional moments. An extension of part of their analysis based on conditional moments is in Gallant, Hansen and Tauchen (1990).

$$\delta_h^k = \log\beta + \frac{1}{2}\left(\gamma^2 \mathrm{Var}\left(\log\frac{c_{h,t+1}}{c_{h,t}}\right) + \mathrm{Var}(\log R_{t+1}^k)\right)$$
$$- \gamma \mathrm{Cov}\left(\log\frac{c_{h,t+1}}{c_{h,t}}, \log R_{t+1}^k\right). \tag{9}$$

Given (8), taking the log of both sides of (5) yields:

$$E[\log R_{t+1}^k] - \gamma E \log\frac{c_{h,t+1}}{c_{h,t}} + \delta_h^k = 0. \tag{10}$$

When a risk-free asset is available, (10) and (9) imply

$$E[\log R_{t+1}^k] - \log R_{t+1}^f = \gamma \mathrm{Cov}\left(\log\frac{c_{h,t+1}}{c_{h,t}}, \log R_{t+1}^k\right)$$
$$- \frac{1}{2}\mathrm{Var}(\log R_{t+1}^k). \tag{11}$$

After some re-arrangement, (11) can be written as:

$$\log E[R_{t+1}^k] - \log R_{t+1}^f = \gamma \mathrm{Corr}\left(\log\frac{c_{h,t+1}}{c_{h,t}}, \log R_{t+1}^k\right)$$
$$\cdot \sqrt{\mathrm{Var}\left(\log\frac{c_{h,t+1}}{c_{h,t}}\right)\mathrm{Var}(\log R_{t+1}^k)}, \tag{12}$$

which is the unconditional version of (4) that Mankiw and Zeldes (1991) use to estimate risk aversion. Since the coefficient of relative risk aversion, γ, is the only parameter to be determined, it can be identified by computing the sample equivalents of the unconditional population moments in (12). This approach is appealing because it relates the preference parameters of interest directly to the time-series properties of consumption growth and asset returns. However, since it relies on a single orthogonality condition to estimate γ, it tends to be rather inefficient. Hence, the usefulness of a conditional version of the first-order condition of the maximization problem: this approach can improve the precision of the estimates by adding information about predictable movements in expected consumption growth rates and asset returns.

Following Hansen and Singleton (1983), we can log-linearize the (conditional) Euler equation in (3) under the assumption of isoelastic preferences and obtain:

$$\log\frac{c_{h,t+1}}{c_{h,t}} = \lambda_{h,t}^k + \gamma^{-1}\log R_{t+1}^k + \varepsilon_{h,t+1}^k. \tag{13}$$

If consumption growth and asset returns are jointly lognormal, $\lambda^k_{h,t}$ is given by:

$$\lambda^k_{h,t} = \frac{1}{\gamma}\log\beta + \frac{1}{2}\left(\gamma\,\mathrm{Var}_t\left(\log\frac{c_{h,t+1}}{c_{h,t}}\right) + \frac{1}{\gamma}\mathrm{Var}_t(\log R^k_{t+1})\right) - \mathrm{Cov}_t\left(\log\frac{c_{h,t+1}}{c_{h,t}}, \log R^k_{t+1}\right), \tag{14}$$

where $\mathrm{Var}_t(.)$ and $\mathrm{Cov}_t(.)$ denote variance and covariance conditional upon the information available at time t. The residual $\varepsilon^k_{h,t+1}$ includes expectation errors and is given by:

$$\varepsilon^k_{h,t+1} = \log\left(\frac{c_{h,t+1}}{c_{h,t}}\right) - E_t\left[\log\left(\frac{c_{h,t+1}}{c_{h,t}}\right)\right] + \frac{1}{\gamma}(\log R^k_{t+1} - E_t[\log R^k_{t+1}]). \tag{15}$$

If there are instruments that are uncorrelated with $\varepsilon^k_{h,t+1}$ and with the innovations to $\lambda^k_{h,t}$, (13) can be estimated using the generalized method of moments. Furthermore, if the model is over-identified, it is possible to test for the predictability of deviations from the Euler equation. Notice that one could estimate risk aversion applying the method of moments directly to the non-linear Euler equation in (3). However, Monte Carlo studies by Vissing-Jørgensen (1999) show that in the presence of measurement error, which is common in micro data, the estimator based on the non-linear Euler equation has poorer properties than estimators based on the log-linearized version in (13). Furthermore, method-of-moment estimators based on non-linear Euler equation may not have a unique solution for risk aversion and, even when they do lead to a unique solution, this has poorer properties than that based on log-linearized Euler equations.[2]

3. The C-CAPM and Limited Stock Market Participation

The first-order conditions for the maximization problem in (1) impose statistical restrictions on the co-movement between an individual's consumption and asset returns. If individuals are identical in preferences and wealth, then in equilibrium they will consume the same amount. In this instance, Equation (3) is valid not only for a single individual, but also for aggregate consumption, and substituting per-capita consumption into the conditions above is justified. However, this degree

[2] Further, Banks et al. (1999) find that even if the hypothesis of joint log-normality of consumption and interest rates does not hold (henceforth, second-order conditional moments turn out significant in the log-linearized Euler equation), the estimate of structural parameters, such as the curvature of preferences, based on (13) is not much affected by the exclusion of conditional second moments.

of homogeneity is clearly unrealistic. In a 1982 paper, Constantinides shows that even if agents are heterogeneous in preferences and wealth, under the assumption of frictionless and complete markets, there exists a utility function for the representative individual[3] that satisfies the above conditions. When there are no frictions and markets are complete, we can construct a "representative" agent, because after trading in complete markets, individuals become homogeneous (at the margin), even though they are initially heterogeneous. It is on this ground that much of the empirical literature on the C-CAPM assumes that population-wide aggregate consumption is the relevant measure to test the predictions of the model. However, if for whatever reason, only a subset of agents are involved in asset trade, individuals will no longer be all homogeneous at the margin and the first-order conditions in (3) and (4) will be valid only for those who are at interior solutions with respect to the holding of the relevant asset. In this instance, it is the aggregate consumption of asset holders that should be used to evaluate the model, and not overall aggregate consumption – which includes both the expenditure of the individuals who satisfy the first-order conditions and the expenditure of those who do not. The inference based on population-wide aggregate data would be valid, and population-wide aggregate expenditure would satisfy the relationship between consumption and asset returns stated above, only if non-asset holders' consumption co-varied with returns in the same way as asset holders' spending. Otherwise, individual-level data providing information on consumption and wealth are needed to distinguish between those for whom the Euler equations in (3) can be expected to hold and those for whom they cannot.

In practice, the distinction between asset holders and non-holders can be carried out in two ways. One is to focus on actual asset holdings at time t, as Mankiw and Zeldes (1991) do, and to classify as shareholders those who hold an amount of shares above a certain threshold. Yet, this approach is likely to bias the analysis of the link between consumption and asset returns conditional upon participation, in the presence of unobservable preference heterogeneity. In fact, when the marginal utility of consumption depends not only on consumption, but also on some unobservable time-varying factor[4] that is correlated with the asset holding decision, grouping households on the basis of actual ownership can yield inconsistent estimates. An alternative procedure has been suggested by Attanasio et al. (2002), who distinguish between shareholders and non-shareholders on the basis of the individual likelihood of holding shares. To be more precise, they compute the probability of share-ownership at time t and label as "predicted shareholders" those households whose predicted probability of owning shares is above a cutoff point.

[3] The coefficient of relative risk aversion of such function is a weighted average of the individual coefficients of relative risk aversion. Hence, it is no larger than that of the most risk-averse agent and no smaller than that of the least risk-averse.

[4] As an example, let individual h instantaneous utility be defined as: $U(c_{h,t}, v_{h,t}) = u(c_{h,t})f(v_{h,t})$, where $v_{h,t}$ denotes unobserved heterogeneity. In this instance, the marginal utility of consumption depends not only on $c_{h,t}$, but also on the value $v_{h,t}$ takes on at time t.

This approach can be thought of as an instrumental variable procedure, where, in the first stage, participation probabilities are determined and, in the second, the estimated probabilities are used to instrument actual participation and study its effects on consumption (see Attanasio et al. (2002) for details).

I adopt Attanasio et al. (2002)'s approach and compute the individual likelihood of shareholding using the results from the estimation of the probability of financial market participation reported in the next Section. In order to account for changes in participation over time, I allow for a variable threshold for the probability of participation. The cutoff point is set equal to the probability that ensures that the portion of predicted shareholders is the same as that of actual shareholders in each of the years covered by the sample. Then, for each couple of adjacent dates, once the group of likely shareholders has been defined on the basis of their probability of ownership in the first of the two periods, I can compute their consumption growth and IMRS between t and $t + 1$.

This approach based on predicted share-ownership brings into the analysis a source of noise related to two sources of possible misclassifications. Some individuals who hold shares can have a predicted probability of ownership below the threshold and hence have their consumption not counted in the measure of the IMRS. Similarly, individuals without shares with predicted probability above the threshold will have their consumption inappropriately included. A Monte Carlo study, reported in Appendix B, examines the implications of this noise for the estimators. The main conclusion is that the misclassification may bias upward the results. Hence, the true risk aversion coefficient can be expected to be even lower than the one I find. Some intuition for this upward bias can be obtained by noticing that including the consumption of some randomly-chosen non-shareholders in the measure of the IMRS and excluding that of some randomly-chosen shareholders may reduce the volatility of the stochastic discount factor and its correlation with stock returns, unless the selection process were correlated with the correlation of consumption growth and returns. This implies that, if anything, the misclassification will bias against finding significant differences between shareholders and the representative agent and my estimates turn out to be conservative.

These classification errors affect also the choice of the instruments used in the GMM estimation. However, since in practice the share ownership equation includes time dummies, these errors are unlikely to be correlated with the aggregate instruments (such as lagged interest rates and so on). Furthermore, as the instruments are lagged two periods or more, this procedure should still yield consistent estimates unless there are reasons to believe that the classification errors exhibit serial correlation.

The availability of individual-level data allows to control for the aggregation process not only by pushing the representative agent assumption to a lower level of aggregation, but also by taking the cross-sectional heterogeneity in consumption explicitly into account. Since the purpose of the analysis is to assess the implications of limited financial market participation *ceteris paribus*, for most of the analysis I will assume a shareholders' representative agent and compare the empirical performance of the model under this assumption with its performance under the alternative (standard) hypothesis of population-wide representative agent. The aggregate marginal rate of substitution is then given by $m_{t+1} = \beta U'(E_h c_{h,t+1})/U'(E_h c_{h,t})$, where E_h denotes the cross-sectional mean taken over the relevant group of households at time t.[5] However, when basing the inference on the log-linearized Euler equation, I take cross-sectional heterogeneity into account by aggregating directly the first-order conditions and focusing on $m_{t+1} = E_h(\beta U'(c_{h,t+1})/U'(c_{h,t}))$. The reason for such choice is to ensure the validity of the instruments in the GMM estimation. In fact, Attanasio and Weber (1993) show that taking logarithms of arithmetic means, which is normally done when using per-capita data and log-linearized Euler equations, induces strong serial correlation in the error, and can lead to rejecting the model due to the correlation of the error with the instrument. Instead, no sign of mis-specification is recorded when using geometric means. Notice that taking cross-sectional heterogeneity in consumption into account and assuming a representative agent within each group would yield similar results if the cross-sectional variance of consumption growth is uncorrelated with asset returns, when (lagged values of) the latter are included in the instrument set.

4. Who Owns Financial Assets in the U.S.?

When deciding how to structure their portfolio, U.S. households face an array of options, ranging from simple bank accounts to highly sophisticated real estate investment trusts. Nevertheless, the portfolios of most American households tend to be very simple and safe. Most households hold a bank account, but relatively few have an investment in stocks and, overall, the ownership of many types of assets is heavily concentrated at the top of the income distribution.

Tables I and II report the patterns of asset ownership in the U.S. according to the U.S. Consumer Expenditure Survey (CEX), the data set that will be used here to assess the empirical performance of the C-CAPM. Although the information on asset ownership and portfolio composition is not very detailed, making the identification of shareholders somewhat imperfect, the CEX has several advantages over other data sets, such as a relatively long and frequent time coverage, a (limited) panel dimension, and information on wealth as well as consumption. As to household portfolios, the CEX provides separate information on only four types of assets: 1. checking, brokerage and other accounts; 2. savings accounts; 3. U.S. savings

[5] The stochastic discount factor based on aggregate data is robust to measurement error in individual consumption, under standard assumptions regarding the latter. Specifically, if it is multiplicative and independent of the true consumption level and asset returns are stationary over time, as the sample size goes to infinite, i.e., as $H_t \to \infty$, measurement error has no effect, since $E_h(c_{h,t}) = E_h(c^*_{h,t}\xi^*_{h,t}) = E_h(c^*_{h,t})E_h(\xi^*_{h,t})$, where $c_{h,t}$, $c^*_{h,t}$ and $\xi^*_{h,t}$ denote observed consumption, true consumption and measurement error, respectively.

Table I. Asset ownership over time

Percentages of households with the possible combinations of assets available in each of the years of the survey and in the sample as a whole. "Risky assets" includes stocks, bonds, mutual funds and other securities.

Year	Percentage of households with:				
	No financial assets	Checking accounts only	Savings accounts only	Checking and savings accounts	Checking/savings accounts and risky assets
1982	0.161	0.105	0.076	0.364	0.292
1983	0.155	0.128	0.074	0.367	0.273
1984	0.163	0.125	0.078	0.348	0.283
1985	0.148	0.135	0.065	0.350	0.298
1986	0.162	0.138	0.058	0.327	0.312
1987	0.151	0.123	0.064	0.329	0.328
1988	0.158	0.140	0.064	0.316	0.320
1989	0.158	0.133	0.053	0.328	0.321
1990	0.137	0.139	0.065	0.343	0.312
1991	0.152	0.135	0.052	0.328	0.328
1992	0.153	0.145	0.057	0.308	0.332
1993	0.156	0.144	0.057	0.296	0.344
1994	0.139	0.158	0.047	0.321	0.328
1995	0.151	0.141	0.053	0.303	0.345
1996[a]	0.155	0.118	0.054	0.321	0.348
Total	0.153	0.135	0.061	0.330	0.316

[a] 1996 refers only to the first quarter.

bonds; 4. stocks, bonds, mutual funds and other securities. "U.S. savings bonds" and "stocks, bonds, mutual funds and other securities" are lumped together and labeled "risky assets", and those households that have non-null holdings are defined as "shareholders". Hence, shareholders will inevitably include many households that do not hold stocks, but only bonds. However, this "imperfect" distinction involves a bias against finding significant differences between shareholders and non-shareholders. As a consequence, it will bias the analysis against, rather than in favor of the prior that the C-CAPM may be an accurate description of shareholders' behavior but should not hold for non-shareholders. Overall, the sample consists of 34,033 households and covers the period from 1982 to 1996. Details on the data and on the criteria of household selection and exclusion are in Appendix A. In the sample, 77, 68 and 32 percent of households hold checking accounts, savings accounts and risky assets, respectively.

Table II. Household characteristics

Socio-demographic characteristics of the set of households without assets, with just checking and/or savings accounts, with checking and/or savings accounts and risky assets, and for the sample as a whole. The sample includes all years covered by the survey. "Risky assets" includes stocks, bonds, mutual funds and other securities. Asterisks (*) denote "fraction of households" with the corresponding characteristics. Income and expenditure are in 1982–1984 U.S. dollars. Standard deviations in parentheses.

	Households with:			Total sample
	No financial assets	Checking/ savings accounts	Checking/savings accounts and risky assets	
Age of head	42.845	43.494	45.706	44.080
Less than high school*	0.497	0.186	0.072	0.198
High school diploma*	0.454	0.587	0.518	0.545
College degree*	0.049	0.227	0.410	0.258
Male head*	0.491	0.665	0.755	0.667
White*	0.631	0.860	0.929	0.846
Black*	0.332	0.097	0.038	0.115
Asian*	0.019	0.018	0.012	0.016
With children*	0.555	0.401	0.428	0.434
Household size	3.104	2.664	2.796	2.774
	(1.912)	(1.561)	(1.405)	(1.582)
Single person*	0.235	0.272	0.186	0.239
North-East*	0.199	0.178	0.209	0.191
Mid-West*	0.216	0.241	0.259	0.243
South*	0.375	0.290	0.272	0.297
West*	0.210	0.291	0.260	0.268
After-tax income	\$11,600	\$22,600	\$33,900	\$24,500
	(10,100)	(15,800)	(20,900)	(18,500)
Monthly consumption	\$480	\$710	\$950	\$750
	(280)	(390)	(520)	(450)
Quarterly consumption growth	0.083	0.100	0.112	0.101
	(0.316)	(0.318)	(0.314)	(0.317)
Quarterly per-capita consumpt. growth	−0.005	0.002	0.002	0.001
	(0.061)	(0.045)	(0.054)	(0.041)
No. of households	5,206	17,933	10,747	34,033

Table I reports the percentages of households with the possible combinations of assets available. I distinguish among households with no assets (15 percent of the sample), with only checking accounts (14 percent), only savings accounts (6 percent), only checking and savings accounts (33 percent), and both risky assets plus checking and/or savings accounts (32 percent). Two features are particularly remarkable: first, many households do not hold any financial asset; second, most hold just few of the available assets. The most significant trend is the increase in the share of households with risky assets, from 28 percent in the early 1980s to 34 percent in the 1990s, which is consistent with the evidence based on the Survey of Consumer Finance, reported by Poterba and Samwick (1995) and Bertaut and Starr-McCluer (2002).

Table II relates portfolio composition to some household characteristics, such as age, education, gender, race, household composition, income and expenditure. With respect to Table I, I have grouped together those holding just checking or savings accounts (riskless asset holders), who represent around 53 percent of the sample. Hence, the table distinguishes between households without assets, households with just riskless assets and households holding also stocks and bonds. Since the differences in terms of demographic characteristics across groups have remained basically unchanged over time (once one controls for the changes occurred in the sample as a whole), the observations for all the years are pooled together. Financial market participation, and especially risky asset ownership, appears to increase with age, education and income, which is consistent with the view that financial market participation is costly. The observed correlation between income and education, on the one hand, and financial market participation, on the other, suggests that such costs may be monetary or may be related to the opportunity costs of investors' time in gathering and processing the information needed to participate to financial markets (see Haliassos and Bertaut (1995) for further discussion). Also, the households participating in financial markets are more likely to be white and headed by a male. Risky asset holding is quite uniformly distributed around the country, whereas no-asset holders are concentrated in the South. In terms of size and composition, households with risky assets are somewhat larger and more likely to have children than those holding just checking and savings accounts, but are smaller and less likely to have children than those holding no assets. The latter are often female-headed, and a larger share are non-white. Table II reports also mean household expenditure on non-durables and services and mean quarterly consumption growth, which, as expected, are lowest among those households with no financial assets and highest among those holding risky assets. The last row reports the rate of growth of consumption of the representative agent, i.e., the household with mean consumption. The quarterly growth rate for the representative holder of risky as well as riskless assets is marginally higher than for the representative riskless asset holder, and substantially higher than for the representative household with no financial assets.

The choice of holding risky assets is investigated further in Table III, which reports the results of the estimation of a bivariate choice (probit) model for risky asset ownership. The individual probabilities of holding risky assets are determined as a function of a polynomial in the age of the household head, of his education, gender, race, marital status, earnings, the presence of children in the household and the area of residence. A set of year dummies is also included. The likelihood of holding financial assets appears to be increasing (although not linearly) in age and education. The dummy for children has a positive sign, but is hardly significant. Single-person, female-headed and non-white households appear to be less likely to hold shares. Finally, the probability of holding shares is highest in the North-East and Mid-West and is strongly increasing in earnings.

5. Empirical Implications of Limited Financial Market Participation

The limited financial market participation hypothesis implies that per-capita consumption growth is a poor proxy for individual consumption growth when pricing assets on the basis of their covariance with the individual intertemporal marginal rate of substitution. Since there is no theoretical reason for non-shareholders to adjust their expenditure in response to predictable changes in share returns, including their consumption in the consumption measure for the C-CAPM will lead to inconsistent estimates of preference parameters and to the rejection of the model's predictions. In this section this proposition is appraised in three ways, using a single data set, providing evidence on the extent to which the anomaly illustrated by Hansen and Jagannathan (1991), Mehra and Prescott (1985) and Hansen and Singleton (1982, 1984) can be attributed to the use of aggregate data that do not account for heterogeneous participation. Rather than implementing a formal test, I replicate these studies and verify whether allowing for limited financial market participation changes their results in a direction that is consistent with the predictions of the theory, which would suggest that part of the anomaly can be accounted for by the use of a "representative agent" assumption rather than a more appropriate "representative stockholding agent" assumption.

The three empirical characterizations of the equity premium anomaly that I consider are: (i) the inconsistency of the representative agent's IMRS with Hansen and Jagannathan bounds; (ii) the equity premium puzzle; (iii) the implausibility of the structural estimates of preference parameters based on aggregate data. For each, the implications of heterogeneous participation are assessed by comparing the evidence based on aggregate data with that based on shareholder data. As discussed earlier, for the empirical analysis I use the information on 'likely shareholders', i.e., on the set of households whose predicted probability of holding risky assets at time t is above a given threshold. The threshold is chosen so as to ensure that the share of predicted shareholders equals that of actual shareholders in each of the years of the sample. As a benchmark I take the representative agent of the CEX sample used in the analysis, whose expenditure is known not to match U.S. aggregate per-

Table III. Probability of holding risky assets

Results of the estimation of a bivariate choice (probit) model for risky asset ownership. The left-hand-side variable is a dummy that is equal to one if the household holds risky assets. "Risky assets" includes stocks, bonds, mutual funds and other securities. Asterisks (*) denote dummy variables. A set of year dummies is also included. Standard errors in parentheses.

	Parameter	Marginal effect
Age	4.952	1.691
	(1.909)	(0.652)
Age^2	−7.799	−2.663
	(4.182)	(1.428)
Age^3	6.187	2.113
	(2.892)	(0.988)
High school diploma*	0.786	0.259
	(0.089)	(0.028)
College degree*	1.177	0.430
	(0.098)	(0.034)
Age × High school diploma	−0.239	−0.082
	(0.163)	(0.056)
Age × College degree	−0.314	−0.107
	(0.188)	(0.064)
Children*	0.024	0.008
	(0.020)	(0.007)
Single-person household*	−0.094	−0.032
	(0.022)	(0.007)
Male*	0.120	0.040
	(0.018)	(0.006)
North-East*	0.040	0.014
	(0.023)	(0.008)
Mid-West*	−0.083	−0.028
	(0.022)	(0.007)
South*	−0.145	−0.049
	(0.023)	(0.007)
Black*	−0.590	−0.172
	(0.030)	(0.007)
Asian*	−0.269	−0.085
	(0.064)	(0.018)
Other non-white*	−0.308	−0.096
	(0.053)	(0.015)
Earnings (thousand $)	0.015	0.005
	(0.000)	(0.000)
Constant	−2.700	–
	(0.281)	
p-value year dummies	0.0000	
No. of observations	34,033	
Pseudo R^2	0.1465	

Table IV. Quarterly returns to the S&P 500 CI and to 3-month T-bills (1982: 1st quarter – 1996: 1st quarter)

Mean and standard deviation of the quarterly return to the S&P500 Composite Index and to 3-month U.S. T-bills and minimum and maximum values achieved by these returns over the period considered.

	Mean	Standard deviation	Min.	Max.
S&P500 CI	0.041	0.074	−0.295	0.267
T-bills	0.016	0.006	0.007	0.032
Equity premium	0.025	0.074	−0.309	0.249

capita expenditure exactly, especially in the levels, as shown in Sabelhaus (1998). As a consequence, comparing the evidence on CEX shareholders with that on the "U.S. aggregate" representative agent would be problematic, since any differences might be due not just to a proper accounting of household heterogeneity, but also to fundamental differences in the data sets. Hence, I compare the evidence on CEX shareholders with that on an agent who features mean CEX consumption expenditure.

5.1. HANSEN AND JAGANNATHAN BOUNDS

Hansen and Jagannathan (1991) derive a lower bound on the volatility of the stochastic discount factors that would be consistent with a given set of asset return data. Their strategy consists in constructing minimum variance random variables that are related to asset returns in the same manner as the IMRS. These random variables are on the mean-standard deviation frontier for the IMRSs and delimit the region of admissible pairs of means and standard deviation for the IMRSs that are consistent with observed asset market data. In Figures 1 and 2 the mean-standard deviation bounds coincide with the solid line.[6] The sequences of crosses and triangles represent the mean (on the horizontal axis) and standard deviation (on the vertical axis) of the (quarterly) IMRS of predicted shareholders and of the (CEX) population-wide representative agent, respectively, for different values of risk aversion. The coefficient of risk aversion starts at 0.5 and increases by steps of 0.5. The quarterly discount rate is set to 2 percent. Figure 1 is based on the (standard) assumption of isoelastic utility, which implies that the quarterly IMRS is defined as $\beta(\bar{c}^g_{t+3}/\bar{c}^g_t)^{-\gamma}$, with $g = (s, \text{all})$. $\bar{c}^g_t$ is "group" mean expenditure on non-durables and services in month t, where the "group" includes predicted shareholders only ($g = s$), or all CEX households ($g = \text{all}$). γ is the Arrow-Pratt coefficient of relative risk aversion. In Figure 2, I consider the case of generalized

[6] The bounds have been computed following Hansen and Jagannathan (1991) using quarterly returns on the S&P500 CI and on U.S. T-bills over the period 1982–1996, reported in Table 4.

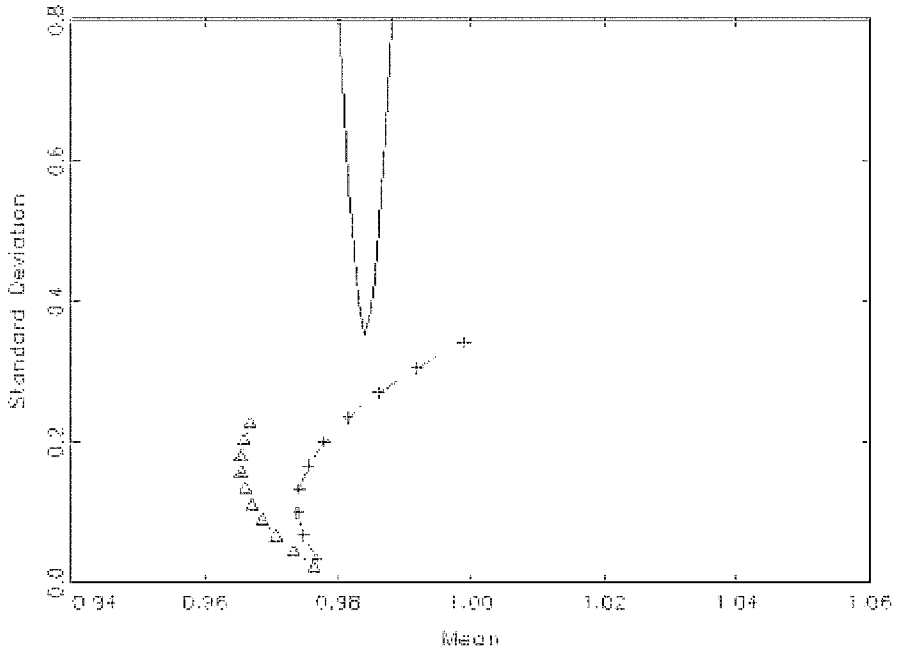

Figure 1. Hansen and Jagannathan bound: the case of isoelastic utility. The cone delimits the region of stochastic discount factors that are consistent with observed asset returns and is based on quarterly returns on the S&P500 CI and T-bills over the period 1982–1996. The symbol + denotes the sequence of IMRSs of predicted shareholders' representative agent. The symbol Δ denotes the sequence of IMRSs of the (CEX) population-wide representative agent. The IMRSs are based on quarterly CEX expenditure on non-durables and services; the discount rate is set to 2 percent; the relative risk aversion coefficient starts at 0.5, increases by steps of 0.5 and runs up to 5.

expected utility as described by Epstein and Zin (1989, 1991). The IMRS is defined as $\left[\beta(\bar{c}^g_{t+3}/\bar{c}^g_t)^{-\rho}\right]^{\frac{1-\gamma}{1-\rho}} R_{t+3}^{\frac{1-\gamma}{1-\rho}-1}$ with $g = (s, \text{all})$.[7] R_{t+3} is the quarterly return on the market portfolio. Following Epstein and Zin (1991), R_{t+3} is set equal to the return on a share index, which, in the case considered, corresponds to the return on the S&P500 CI. ρ is the inverse of the intertemporal elasticity of substitution. In the literature, there is no consensus as to the value of the elasticity of substitution. Hall (1988) supports the view of an elasticity below one; other papers, such as Mankiw et al. (1985), obtain higher values that are generally above one. Hence, I set ρ equal to 5 and 0.5 and consider both the case of an elasticity of 0.2 and that of an elasticity of 2.

[7] In the case of generalized expected utility, to justify aggregation over consumers and to apply the model to data on per-capita consumption, one needs to adopt the common (although restrictive assumption) of identical and homothetic preferences. Wealth can vary across consumers. See Epstein and Zin (1991) for details.

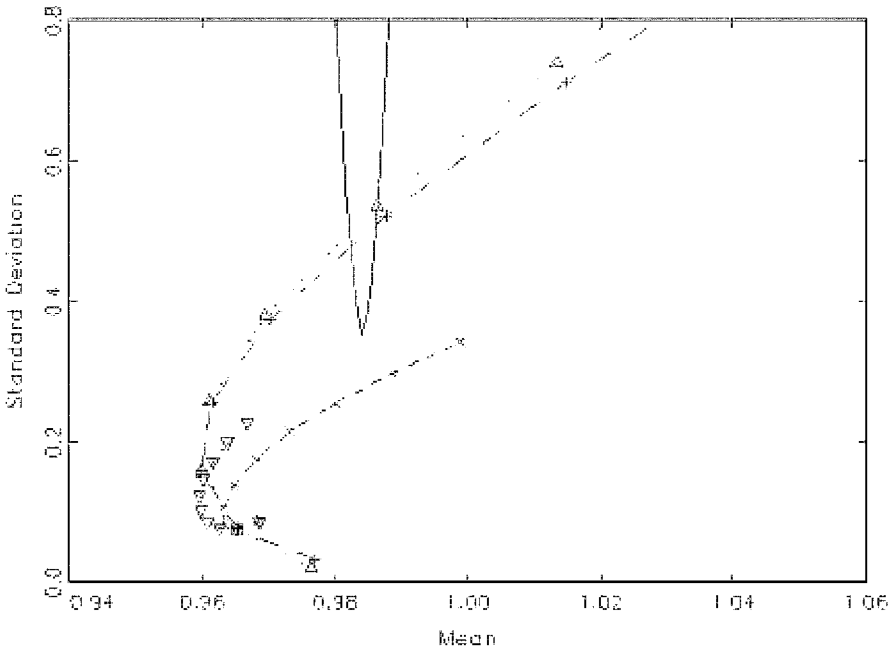

Figure 2. Hansen and Jagannathan bound: the case of Epstein and Zin utility. The cone delimits the region of stochastic discount factors that are consistent with observed asset returns and is based on quarterly returns on the S&P500 CI and T-bills over the period 1982–1996. The symbols + and × denote the sequences of IMRSs of predicted shareholders' representative agent with elasticity of substitution equal to 2 and 0.2, respectively. The symbols Δ and ∇ denote the sequences of IMRSs of the (CEX) population-wide representative agent with elasticity of substitution equal to 2 and 0.2, respectively. The IMRSs are based on quarterly CEX expenditure on non-durables and services; the discount rate is set to 2 percent; the relative risk aversion coefficient starts at 0.5 and increases by steps of 0.5.

The criterion that will be used to appraise the evidence is based on a linear combination of the coefficient of relative risk aversion and the Euclidean distance of the IMRS from the Hansen and Jagannathan bound, such as $C = d + \omega\gamma$, where γ denotes the coefficient of relative risk aversion, d the distance of the IMRS computed at that γ and ω their relative weight.[8] For any value assigned to ω, and for γ smaller than 6, the IMRS of the representative shareholder (crosses) minimizes the criterion. For shareholders, the Euclidean distance of the IMRS from the bound is minimized for a risk aversion around 4. For the population-wide representative agent the minimum is achieved at a much higher value. The Euclidean distance from the frontier of the mean-variance pair of the IMRS based on predicted shareholders' expenditure is smaller, due to the smaller variance of

[8] Computing confidence bands for the IMRS series is not straightforward. However, since the sample of shareholders is less than a third of the overall sample of households, one can expect the confidence intervals for the IMRS based on shareholder data to be much wider than those for the IMRS based on the expenditure of all households.

non-shareholder consumption growth, which lowers the variance of aggregate consumption growth. Hence, for the representative agent, to get close to the bound a larger value of γ is needed. This suggests that distinguishing between shareholders and non-shareholders may be important for improving the predictive potential of the model.

Figure 2 considers the implications of Epstein and Zin utility for different values of the elasticity of substitution. When the elasticity of substitution is set to 0.2, the IMRS of predicted shareholders (denoted by the symbol $\times$) minimizes the criterion combining degree of risk aversion and Euclidean distance from the bound for values of risk aversion below 6, for any ω. With an elasticity of substitution equal to 2, the IMRS for predicted shareholders (denoted by the symbol +) falls in the region delimited by the bound for a risk aversion of 3.5. Interestingly, the IMRS for the representative agent (denoted by the symbol Δ) is also consistent with the bound for similar levels of risk aversion. This result is at odds with the hypothesis that accounting for heterogeneity in participation is crucial for pricing assets correctly. Yet, it does provide evidence favorable to the theories that assign a role to preference non-separabilities in explaining the empirical failure of the C-CAPM.

5.2. CALIBRATION OF THE RISK AVERSION COEFFICIENT

The focus of the analysis based on the calibration of the equity premium implied by the data is on the correlation of consumption with asset returns and on its role in rationalizing the empirical evidence. The model is estimated both in its non-linear form – which allows us to examine the implications of alternative specifications of the felicity function – and after log-linearizing it, in the spirit of Mankiw and Zeldes (1991).

5.2.1. *The Non-linear Model*

In this section, I basically re-do Mehra and Prescott (1985) in that I calibrate the model parameters to match asset market data. Table V reports the minimum value of the risk aversion parameter that is consistent with observed behavior, i.e., that makes $\bar{e}^g = (T-3)^{-1}\sum_t e^g_{t+3}$ statistically indistinguishable from zero, given the time-discount factor, where $\bar{e}^g_{t+3} = [\beta U'(\bar{c}^g_{t+3})/U'(\bar{c}^g_t)](R^{S\&P500}_{t+3} - R^{T\text{-}bills}_{t+3})$, with $g = (s, \text{all})$. $U'(\,)$ is the marginal utility of consumption. $\bar{c}^g_t$ is "group" mean expenditure on non-durables and services, where the group includes predicted shareholders only ($g = s$) and all CEX households (g = all). β denotes the time-discount factor.

Table V shows the results of the calibration under the hypotheses of isoelastic preferences and of Epstein and Zin utility. When preferences are isoelastic, focusing on shareholders lowers the coefficient of relative risk aversion that reconciles observed expenditure behavior with asset returns. In fact, given a quarterly discount rate of 2 percent, γ drops from 9 for the CEX representative agent to 7 for the representative agent of (predicted) shareholders. The evidence on Epstein

Table V. Calibration of the risk aversion coefficient (the non-linear model)

Calibration of the coefficient of risk aversion based on the non-linear C-CAPM, under alternative preference specifications. The table reports the minimum value of the risk aversion parameter (γ) that makes $\bar{e} = (T-3)^{-1} \sum_t [\beta U'(\bar{c}_{t+3})/U'(\bar{c}_t)] (R_{t+3}^{S\&P500} - R_{t+3}^{T\text{-}bills})$ statistically indistinguishable from zero. $U'(\cdot)$, $\bar{c}_t$ and β denote the marginal utility of consumption, average consumption and the time-discount factor, respectively. The quarterly discount rate is set equal to 2 percent. ρ is the inverse of the elasticity of substitution (*e.s.*). Standard errors in parentheses. The standard errors have been corrected to account for the MA(6) structure of the residual. The problem of non-positive definite variance-covariance matrix in finite samples has been taken care of by using a set of weights, as in Newey and West (1987).

	Isoelastic utility		Epstein and Zin utility			
			$\rho = 0.5$ (*e.s.* = 2)		$\rho = 5$ (*e.s.* = 0.2)	
	e	γ	e	γ	e	γ
Predicted	0.023	7	2.915	8.5	300.343	59
shareholders	(0.012)		(1.529)		(152.875)	
Whole sample	0.022	9	2.592	8	56.578	62
	(0.011)		(1.350)		(29.060)	

and Zin utility is consistent with that emerging from the Hansen and Jagannathan-type analysis in Figure 2. When the elasticity of substitution is set equal to 2, predicted shareholders are largely indistinguishable from non-shareholders and the C-CAPM turns out to be consistent with the empirical evidence for a relative risk aversion around 8, with no need to account for heterogeneous participation. When the elasticity of substitution is less than 1, the estimated risk aversion for predicted shareholders is lower than for the representative agent but turns out to be around 60.

5.2.2. *The Log-linear Model*

Table VI shows the risk aversion estimates based on the calibration of the unconditional log-linearized C-CAPM of equation (12), whose sample analogue is:

$$\gamma^g = \frac{\log \overline{R^{S\&P500}} - \log \overline{R^{T\text{-}bills}}}{\widehat{\text{Corr}}(\overline{\Delta \log c_{t+3}^g}, \log R_{t+3}^{S\&P500}).\sqrt{\widehat{\text{Var}}(\overline{\Delta \log c_{t+3}^g})\widehat{\text{Var}}(\log R_{t+3}^{S\&P500})}},$$

$$g = (s, \text{all}), \tag{16}$$

where $\overline{R^i} = (T-3)^{-1} \sum_{t+1}^{T-3} R_{t+3}^i$, $i = (S\&P500,\ T\text{-}bills)$ and $\widehat{\text{Corr}}(.)$ and $\widehat{\text{Var}}(.)$ denote the sample correlation and variance; $\overline{\Delta \log c_{t+3}^g} =$

Table VI. Calibration of the relative risk aversion coefficient (the log-linear model)

Calibration of the coefficient of risk aversion based on the unconditional log-linearized C-CAPM. The coefficient of risk aversion (γ) is determined using the following relationship:

$$\gamma = \left(\log \overline{R^{S\&P500}} - \log \overline{R^{T\text{-}bills}}\right) \times \left(\widehat{\mathrm{Corr}}(\overline{\Delta \log c_{t+3}}, \log R_{t+3}^{S\&P500}) \cdot \sqrt{\widehat{\mathrm{Var}}(\overline{\Delta \log c_{t+3}})\widehat{\mathrm{Var}}(\log R_{t+1}^{S\&P500})}\right)^{-1},$$

$$\log \overline{R^{S\&P500}} - \log \overline{R^{T\text{-}bills}} = 0.0250.$$

Column (1) and (2) display the correlation between consumption growth and the return on the S&P500 CI and the variance of consumption growth, respectively.

	(1)	(2)	γ
Predicted shareholders	0.143	0.054	44
Whole sample	0.130	0.036	72
Predicted non-shareholders	*0.121*	*0.030*	–

$(H_t^g)^{-1} \sum_{h \in H_t^g} \Delta \log c_{h,t+3}^g$ with H_t^g denoting the set of (predicted) shareholders ($g = s$) and of CEX households (g = all) interviewed both at date t and at $t + 3$. The rate of growth of consumption is based on the geometric mean of individual consumption ratios. The evidence based on arithmetic means of individual consumption observations is qualitatively similar.

The table reports the correlation between consumption growth and the (log) return on the S&P500 CI, the standard deviation of consumption growth and the corresponding relative risk aversion coefficient. When aggregating across all CEX households, the implied relative risk aversion coefficient is 72. When allowing for heterogeneity in financial market participation and singling out shareholders, it drops to 44. Calibrating the model using arithmetic means, i.e., assuming that the population-wide and the shareholders' representative agents consume mean expenditure, yields relative risk aversion coefficients equal to 60 and 48, respectively.[9]

The difference between the estimate based on shareholder data and that based on all households stems from differences in the pattern of consumption *vis à vis* asset returns between shareholders and non-shareholders: (predicted) shareholders' consumption growth is more volatile than non-shareholders' and the correlation with risky asset returns is much stronger, as the table shows. This is why including non-shareholders' expenditure in the consumption measure used to price shares will bias the inference.

9 Since the estimates are the result of a back-of-the-envelope calculation, standard errors are not reported, but as they are based on only one moment restriction, they can be expected to be quite large.

By comparison with Mankiw and Zeldes (1991), the results obtained in this section differ slightly from a quantitative point of view, but not qualitatively. In fact, Mankiw and Zeldes find somewhat larger differences in the time-series properties of consumption of shareholders and non-shareholders, but estimate the coefficient of risk aversion for the whole PSID sample and for shareholders at 100 and 35, respectively.

5.3. STRUCTURAL ESTIMATION OF THE PREFERENCE PARAMETERS

The results of the structural estimation of the preference parameters based on the log-linearized C-CAPM in (13) are reported in Table VII. The estimated relation is the following:

$$\overline{\Delta \log c^{g}_{t+3}} = \lambda^{i} + \gamma^{-1} \log R^{i}_{t+3} + \varepsilon^{g,i}_{t+3},\ i = (S\&P500,\ T\text{-}bills);\ g = (s, \text{all}). \tag{17}$$

The left-hand-side variable of (17) is the rate of growth of consumption of predicted shareholders ($g = s$) or of CEX shareholders and non-shareholders (g = all), computed as geometric mean of individual consumption ratios. (17) is estimated for stock returns and for T-bills, separately at first and then jointly to improve the efficiency by exploiting the cross-equation correlation of the residuals arising from the correlation of the expectational errors.

The estimation of the parameter of interest γ is based on the linear generalized method of moments (GMM) and relies on the orthogonality of the residual to the information on the changes of households' expectations of consumption and returns. Hence, it is more efficient than the estimation based on the calibration of the model, which relies on a single restriction. The choice and timing of the instruments must account for the endogeneity of asset returns with respect to the expectational components of the error, as given in (15), and for the quarterly frequency of the left-hand-side variable. Moreover, if the variance and covariance terms in (14) are not constant, their stochastic components enter the residual, which does not cause problems as long as they are uncorrelated with the instruments. Further, the residual of the log-linearized Euler equation turns out to be an MA(6) because of the overlapping of the quarterly consumption growth rates (due to the monthly frequency of interviews) and because each household contributes up to 4 observations on expenditure.[10] Hence, as instruments, I use the average quarterly

[10] The residual of the log-linearized Euler equations results from two components: an expectation error and a measurement error. The former is white noise, or at most an MA(1) due to presence of time aggregation effects, at the household level. However, since households are interviewed every month of the year, the observations on quarterly consumption growth overlap and the individual expectation components in $\varepsilon_{h,t+3}$ can be expected to be correlated with the expectation components in $\varepsilon_{h,t+2}$ and $\varepsilon_{h,t+1}$, yielding an MA(2). The measurement error generates an MA(6) component and is due to the fact that, in the absence of attrition, the household that enters the survey at t contributes to the IMRS dated $t + 3$, $t + 6$ and $t + 9$. Hence, the resulting residual of (17) is an MA(6).

Table VII. GMM estimation of the Euler equation

Generalized method of moment estimation based on the log-linearized C-CAPM for stock returns and T-bills, considered separately in the first and second set of rows and jointly in the last set. Constant[a] and Constant[b] refer to the constant terms in the equations with the return on the S&P500 CI and on T-bills, respectively, when they are estimated jointly. The set of instruments include: average consumption growth lagged three and four quarters; the returns on T-bills and on the S&P500 CI and the rate of inflation lagged one to four quarters; and the risk and term spreads lagged one quarter. Standard errors in parentheses. The standard errors have been corrected to account for the MA(6) structure of the residual and a set of weights as in Newey and West (1987) have been used to take care of the problem of non-positive definite variance-covariance matrix in finite samples. The degrees of freedom to the Sargan test of overidentifying restrictions are 15 in the single equation model and 31 in the two-equation model.

		Predicted shareholders	Whole sample
S&P500	γ^{-1}	0.120	0.081
		(0.058)	(0.072)
	Constant	−0.001	−0.003
		(0.003)	(0.004)
	Sargan test	13.660	13.200
	(*p-value*)	(0.552)	(0.587)
T-bills	γ^{-1}	0.740	0.192
		(0.325)	(0.365)
	Constant	−0.007	−0.003
		(0.006)	(0.007)
	Sargan test	16.237	9.688
	(*p-value*)	(0.367)	(0.839)
S&P500	γ^{-1}	0.133	0.091
and T-bills		(0.057)	(0.071)
	Constant[a]	−0.001	−0.004
		(0.003)	(0.004)
	Constant[b]	0.003	−0.002
		(0.002)	(0.002)
	Sargan test	29.806	22.921
	(*p-value*)	(0.527)	(0.852)

rate of growth of consumption (lagged three and four quarters), the quarterly returns on the S&P500 CI and on T-bills (lagged one to four quarters), the rate of inflation (lagged one to four quarters) and the risk and term spreads (lagged one quarter). The latter are computed as differences between the yield on an index of BAA-rated corporate bonds and the yield on an index of AAA-rated corporate bonds (risk spread) and between the yield on an index of long-term U.S. government bonds and the yield on three-month T-bills (term spread). Lagged stock and bond returns are used in many previous studies as instruments. The use of risk and bond horizon premiums is motivated by the findings of Fama and French (1989) that they have strong predictive power for share returns. For evidence on the predictability of asset returns using the other instruments, see, for example, Attanasio and Weber (1995). Other variables in the information set can be expected to be orthogonal to the error term of the Euler equation. However, it is a well-known result that the properties of the estimator can deteriorate if weak instruments are included. First-stage regressions are at the basis of the parsimonious list of over-identifying instruments used here.

Table VII reports the results from the estimation. For the whole sample, the estimates imply a high and imprecisely estimated value of the coefficient of risk aversion, when this is estimated both from the Euler equation for shares and from the equations for shares and T-bills taken jointly. However, the test of over-identifying restrictions does not reject the null, which considering the coefficient estimates may be due to somewhat weak instruments. For the group of predicted shareholders, the point estimates of γ^{-1} are very precise and imply plausible values for the risk aversion coefficient, which turns out equal to 8.3 when using the S&P500 CI equation only and just a bit lower at 7.5 when imposing the cross-equation restriction. Further, the test of over-identifying restrictions never rejects the null. Compared to those in Attanasio et al. (2002), based on the U.K. Family Expenditure Survey, these GMM point estimates of the preference parameters are somewhat larger, but have broadly consistent implications.[11]

6. Conclusions

This paper examines the implications of limited financial market participation for the empirical performance of the Consumption-based Capital Asset Pricing Model (C-CAPM). Traditionally, the C-CAPM has not performed well empirically. The objective of this paper is to assess to what extent the equity premium anomaly illustrated by Hansen and Jagannathan (1991), Mehra and Prescott (1985) and Hansen and Singleton (1982, 1984) can be attributed to the use of aggregate data, hence to the inclusion of non-shareholders' expenditure in the consumption measure used to price assets. I start with Hansen and Jagannathan's framework and compare the IMRSs of shareholders and those of the CEX population-wide representative agent with the Hansen and Jagannathan bound. Hansen and Jagannathan's framework illustrates the role of the IMRS volatility for the C-CAPM. This approach is non-parametric and applies to a rich class of pricing models. It can be thought of as complementary to those based on more traditional tests of overidentifying restrictions, whose success is mixed since they do not provide any indication as to how the model should be changed in order to correct or improve its specification.

[11] Also in Attanasio et al. (2002) the test of overidentifying restrictions fails to reject the null for both the sample of predicted shareholders and the whole sample of households.

Next, I calibrate Mehra and Prescott equity premium and estimate the coefficient of relative risk aversion for shareholders and for the population-wide representative agent, exploiting the unconditional restrictions implied by the model. The calibration allows to appraise the differences in the time-series properties of the consumption profiles of the two groups and their implications for the estimation of preference parameters. Last, I take a structural approach. I use the conditional distribution of the changes in expectations of consumption and asset returns to estimate efficiently the parameters of interest and test for the overall suitability of the C-CAPM in describing the data in the spirit of Hansen and Singleton (1982, 1984). Although the three approaches yield somewhat different point estimates of the preference parameters conditional on participation, all three estimates differ from the corresponding unconditional one in the same direction, pointing out that at least part of the equity premium anomaly can be accounted for by the use of a "representative agent" assumption rather than a more appropriate "representative stockholding agent" assumption.

Appendix A: The U.S. Consumer Expenditure Survey

The data used for the analysis are taken from the U.S. Consumer Expenditure Survey (CEX), which is a representative sample of the U.S. population, conducted by the U.S. Bureau of Labor Statistics. The survey is a rotating panel in which interviews occur continuously throughout the year, with each consumer unit being interviewed every three months over a twelve-month period, apart from attrition. About 4,500 households are interviewed each quarter, more or less evenly spread over the three months: 80 percent are re-interviewed after three months, whereas the remaining 20 percent are dropped and a new group is added. So approximately one fifth of the units interviewed each month are new to the survey. This rotating procedure is designed to improve the overall efficiency of the survey and to attenuate attrition.

The sample used cover the period from 1982 to 1996, first quarter.[12] I exclude those households with incomplete income responses and those living in rural areas or in university housing. In addition, I exclude those whose head was under 21 or over 75 years old, those that do not participate in the last interview when the information on financial wealth is collected (about 30 percent of the initial sample), and those whose financial supplement contains invalid blanks in the stocks of assets held (around 20 percent of respondents). Finally, I drop those households with

[12] In 1986, the sample design and the household identification numbers were changed. 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. Since the distinction between shareholders and non-shareholders is based on information collected during the last interview only, those households whose first interview was in the third or fourth quarter of 1985 are excluded. For similar reasons, those whose first interview is after June 1995 are also excluded. Thus, the sample used consists of households whose first interview was between 1982:1 and 1985:6 or between 1986:2 and 1995:6.

no consecutive interviews or reporting zero food expenditure and those whose total consumption on non-durables and services changes by more than 300 percent between contacts. 147 households that reported holding only risky assets are also ignored, as most report negligible stockholdings which might reflect coding errors. As so restricted, the sample consists of 121,777 observations on 34,033 households.

Each quarterly interview collects household monthly expenditure for the previous three months. For my analysis, however, I use only the information for the month preceding that of interview. In practice, many households report to spend exactly the same amount in each of the three months preceding that of interview. This is likely to be a consequence of the fact that this information is based on recall (and not, for example, on diaries) and households are dividing by three their quarterly expenditure. Thus, using all three observations would not add much information. Monthly consumption is seasonally adjusted using twelve month dummies and 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). Given the timing of the observations on expenditure, which refer to the month preceding the one when the interview takes place, each household contributes at most four observations on monthly consumption at quarterly intervals. Overall, the sample spans over 171 months, with consumption observations available from 1981:12 to 1996:2.

The last column of Table II reports descriptive statistics for the sample as whole. All figures are in 1982–1984 U.S. dollars. Mean annual household after-tax income is \$24,500; mean monthly consumption is \$750. Mean quarterly consumption growth is 0.10. The average age of the head is 44. 20 percent of heads do not have a high school diploma; 55 percent have a high school diploma; 26 percent have a college degree. 67 percent of household heads are male; 85 percent are white; 12 percent are black and 24 percent of households are single persons. 43 percent of households have children. 19 percent live in the North-East, 24 percent in the Mid-West, 30 percent in the South and 27 percent in the West.

Table IV summarizes the evidence on the real quarterly returns used in the analysis. Risky returns are measured by the return on the Standard & Poor's 500 Composite Index (S&P500 CI). Unsurprisingly, over the period 1982-1996, the average risky return is substantially higher and more volatile than the return on 3-month Treasury Bills (T-bills), which is taken to be the riskless asset. The mean quarterly premium is more than 2.5 percent.

Appendix B: Monte Carlo Study of the Effects of Using Predicted Shareholders Rather than Actual Shareholders

In the analysis of the link between consumption and asset returns conditional upon financial market participation, following Attanasio et al. (2002) I distinguish between shareholders and non-shareholders based on the individual likelihood of holding shares in order to avoid any bias that may arise in the presence of unobservable preference heterogeneity. Two types of possible mis-classification can occur. Individuals who hold shares can have a predicted probability of ownership that falls below the threshold and be classified as 'predicted non-shareholders'. Individuals who do not hold shares can have a predicted probability above the threshold and be classified as 'predicted shareholders'. The misclassification is random; nevertheless it is a source of noise and may affect the analysis. The purpose of this Appendix is to assess the implications of this noise for the risk aversion estimators ($\hat{\gamma}$) used in the paper.

The Monte Carlo simulation consists of repeating the following steps 5,000 times in order to characterize the distribution of the estimator of interest.

Step 1. Draw a time series for each of the following variables: shareholders' consumption growth, non-shareholders' consumption growth and equity premium. These series are jointly log-normal and have the same time-series properties (mean and variance-covariance matrix) of my CEX series for the 'representative' predicted shareholder and non-shareholder, i.e., for the agents with mean predicted shareholders' and non-shareholders' consumption.

Step 2. For each t, H^g idiosyncratic shocks $\varepsilon^g_{h,t+1}$ are then drawn for shareholders ($g = s$) and non-shareholders ($g = ns$). These are assumed to be lognormal, $ln(\varepsilon^g_{h,t}) \sim N(0, \sigma^2_{ln(\varepsilon^g)})$ and independent over time, from each other and from consumption growth and asset returns. These shocks represent cross-sectional variation in consumption. The value of $\sigma^2_{ln(\varepsilon^g)}$ is based on the data sample that I use. The cross-sectional standard deviation of consumption growth for predicted shareholders (non-shareholders) is equal to 0.51 (0.48). Since the standard deviation of consumption growth for the representative predicted shareholder (non-shareholders) is 0.07 (0.04), the standard deviation of $\Delta ln(\varepsilon^g_{h,t+1})$ turns out equal to $(0.51^2–0.07^2)^{1/2} = 0.505$ for shareholders and $(0.48^2–0.042)^{1/2} = 0.479$ for non-shareholders. With independent shocks over time, $\sigma^2_{\Delta ln(\varepsilon^g)} = 2\sigma^2_{ln(\varepsilon^g)}$. Hence, for shareholders and non-shareholders, $\sigma^2_{ln(\varepsilon^g)}$ are 0.357^2 and 0.338^2, respectively. Since this is only a rough calibration, the sensitivity to various values of this parameter are considered.

Step 3. At each t, randomly replace a fraction of shareholders with non-shareholders and compute the rate of growth of consumption for this group of noisily predicted shareholders'. In the data, about 35 percent of predicted shareholders do not actually hold shares and about 17 percent of actual non-shareholders are predicted to hold shares. These are the rates of replacement that I use. However, since the extent of the noise due to misclassification depends on these rates

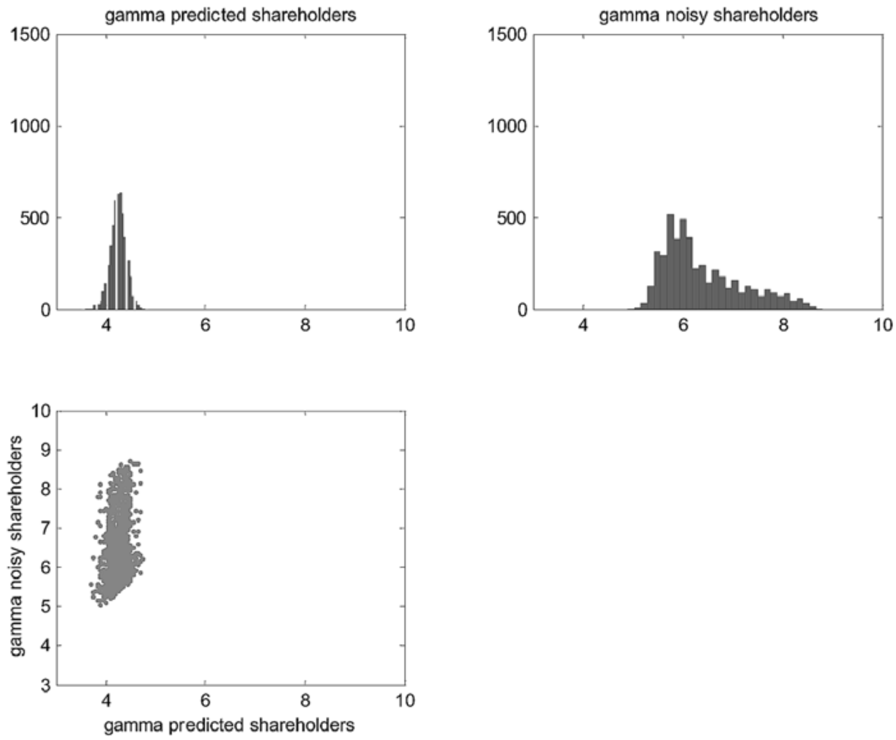

Figure 3. Monte Carlo simulation of the properties of risk aversion estimators with misclassified shareholders and non-shareholders: *Hansen and Jagannathan bound* (the case of isoelastic utility). Distribution of the coefficient of risk aversion based on 5,000 Monte Carlo iterations. The coefficients are obtained by minimizing the distance of the IMRS sequences from the Hansen and Jagannathan bound. The standard deviations of the idiosyncratic shocks to shareholders' and non-shareholders' consumption growth are set to 0.505 and 0.478, respectively. The sample of noisy shareholders is obtained by replacing 35 percent of predicted shareholders with non-shareholders. The bottom panel plots against each other 1,500 randomly chosen points corresponding to 1,500 of the 5,000 simulations.

and since the focus of this appendix is on the implications for the estimators of structural parameters of such noise, I consider also other values.

Step 4. Calculate the value of each estimator using the sample of 'predicted shareholders' from Step 1 and using the sample of 'noisily-predicted shareholders' from Step 3.

Figures 3 through 5 and Table VIII report the results of the simulations for three of the risk aversion estimators ($\hat{\gamma}$) used in the paper. All are based on the Euler equation for the excess return on stocks over T-bills. For simplicity I do not consider the instrumental variable estimator, but the basic points carry over to this case. The figures plot the distributions of the risk aversion estimators for $\sigma_{\Delta ln(\varepsilon^{g}_{h,t+1})}$ equal to 0.505 and to 0.479 for predicted shareholders and non-shareholders, respectively. The top left-hand panel of each figure shows the distribution of the

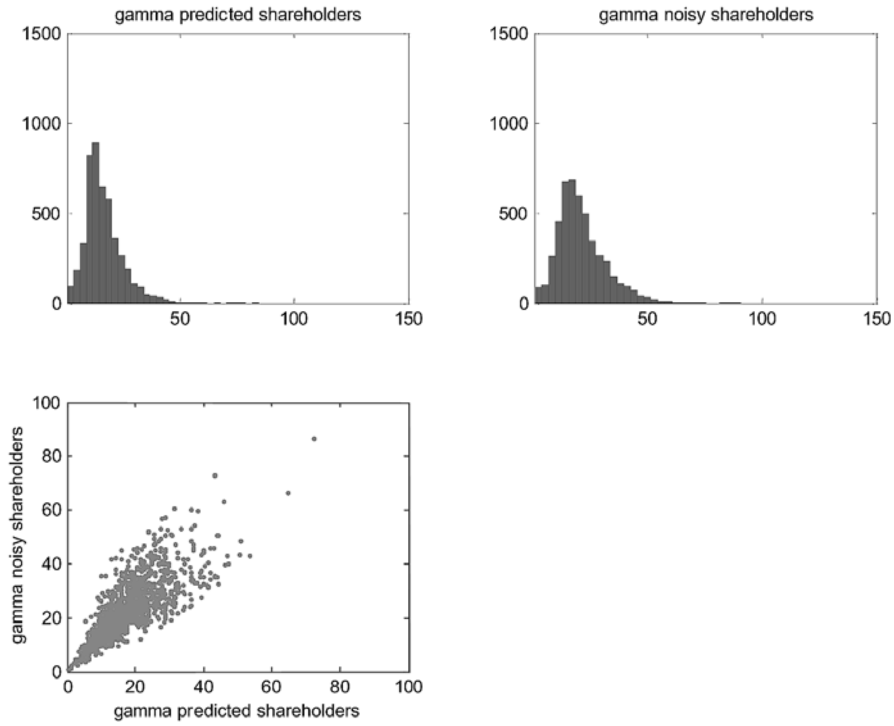

Figure 4. Monte Carlo simulation of the properties of risk aversion estimators with misclassified shareholders and non-shareholders: *calibration of the non-linear model*. Distribution of the coefficient of risk aversion based on 5,000 Monte Carlo iterations. The coefficients are obtained by calibrating the non-linear C-CAPM. The standard deviations of the idiosyncratic shocks to shareholders' and non-shareholders' consumption growth are set to 0.505 and 0.478, respectively. The sample of noisy shareholders is obtained by replacing 35 percent of predicted shareholders with non-shareholders. The bottom panel plots against each other 1,500 randomly chosen points corresponding to 1,500 of the 5,000 simulations.

risk aversion estimator based on the sample of predicted shareholders; the top right-hand panel shows the distribution of the estimator when some noise due to misclassification is thrown in, by randomly replacing 35 percent of predicted shareholders with non-shareholders; the bottom panel plots against each other about 1,500 randomly-chosen points corresponding to 1,500 of the 5,000 Monte Carlo iterations. If the misclassification had little effect, then the distributions should be close-to-identical and the points in the bottom panel should lie close to the 45-degree line. The table provides detailed statistics for various values of the $\sigma_{\Delta ln(\varepsilon^g_{h,t+1})}$ and of the degree of misclassification.

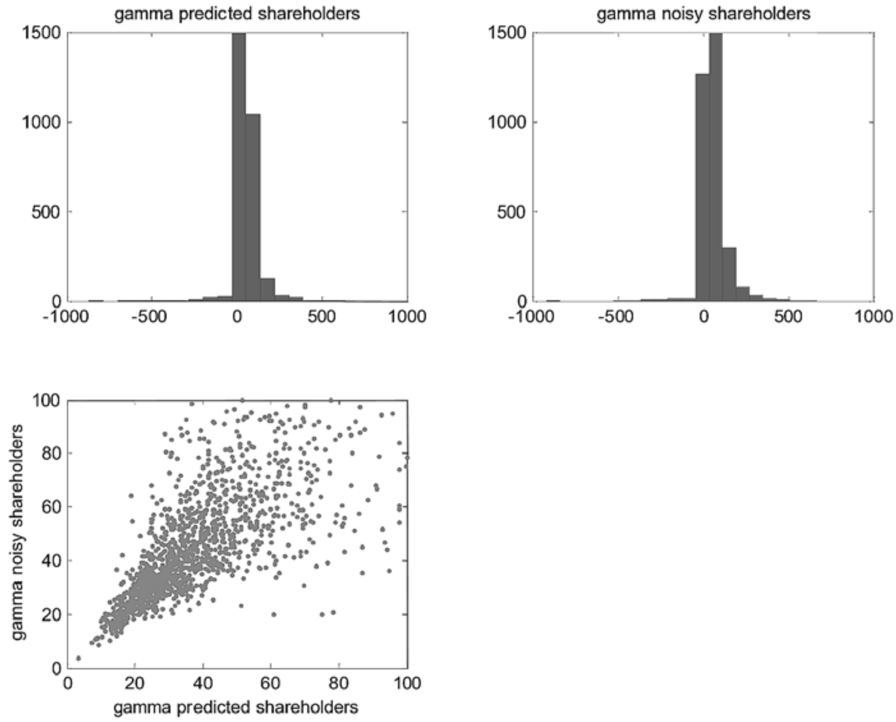

Figure 5. Monte Carlo simulation of the properties of risk aversion estimators with misclassified shareholders and non-shareholders: *calibration of the log-linear model*. Distribution of the coefficient of risk aversion based on 5,000 Monte Carlo iterations. The coefficients are obtained by calibrating the unconditional log-linearized C-CAPM. The standard deviations of the idiosyncratic shocks to shareholders' and non-shareholders' consumption growth are set to 0.505 and 0.478, respectively. The sample of noisy shareholders is obtained by replacing 35 percent of predicted shareholders with non-shareholders. The bottom panel plots against each other 1,500 randomly chosen points corresponding to 1,500 of the 5,000 simulations.

1. IMPLICATIONS OF MISCLASSIFICATION FOR HANSEN AND JAGANNATHAN BOUND-BASED ANALYSIS

At each simulation round, I compute Hansen and Jagannathan bounds and the IMRSs of predicted shareholders and of the noisily-predicted shareholders, which are given by:

$$\frac{1}{T}\sum_t \left(\frac{\bar{c}_{t+1}}{\bar{c}_t}\right)^{-\gamma}, \tag{18}$$

for different values of γ and plot them against the corresponding bound. The parameter γ starts at 0.5 and increases by steps of 0.05 until 10. Figure 3 plots the distribution of the $\hat{\gamma}_s$ that minimize the Euclidean distance of the IMRS sequence from the bound for predicted shareholders and for the noisily-predicted shareholders. Table VIII reports their mean and standard deviation and the mean and standard deviation of their difference, for several values of the $\sigma_{\Delta ln(\varepsilon^{g}_{h,t+1})}$ and of the rate of misclassification. The two estimators turn out to be different: the $\hat{\gamma}$ based on the sample of noisily-predicted shareholders is systematically larger than that based on the sample of predicted shareholders. The difference is increasing in the degree of misclassification, whereas the impact of differences between $\sigma_{\Delta ln(\varepsilon^{s}_{h,t+1})}$ and $\sigma_{\Delta ln(\varepsilon^{ns}_{h,t+1})}$ is non-linear. Overall, the analysis implies that the random misclassification may be a source of upward bias. Hence, my estimates of preference parameters are "conservatively small" and the true γ can be expected to be even lower than the one that I find.

2. IMPLICATIONS OF MISCLASSIFICATION FOR THE CALIBRATION OF THE COEFFICIENT OF RISK AVERSION: THE NON-LINEAR MODEL

The next estimator is based on the calibration of model parameters to match asset market data in the spirit of Mehra and Prescott (1985). $\hat{\gamma}$ solves:

$$\frac{1}{T}\sum_{t}\left(\frac{\bar{c}_{t+1}}{\bar{c}_{t}}\right)^{-\hat{\gamma}}(R^{S\&P500}_{t+1} - R^{T\text{-}bill}_{t+1}) = 0. \tag{19}$$

This estimator is based on the non-linear Euler equation and is calculated using grid search, over the interval 0 and 150 with 301 grid points, as the grid point which sets the moment equal to zero. Figure 4 reports the distribution of the estimators for predicted shareholders and for the noisily-predicted shareholders and summary statistics are in Table VIII: the 'noisy estimator' turns out to be larger on average than the estimator based on predicted shareholders.

3. IMPLICATIONS OF MISCLASSIFICATION FOR THE CALIBRATION OF THE COEFFICIENT OF RISK AVERSION: THE LOG-LINEAR MODEL

The last estimator is based on the calibration of the unconditional log-linearized C-CAPM and $\hat{\gamma}$ is given by:

$$\hat{\gamma} = \frac{\log \overline{R^{S\&P500}} - \log \overline{R^{T\text{-}bills}}}{\hat{\mathrm{Corr}}(\overline{\Delta \log c^{g}_{t+1}}, \log R^{S\&P500}_{t+1}).\sqrt{\hat{\mathrm{Var}}(\overline{\Delta \log c^{g}_{t+1}})\hat{\mathrm{Var}}(\log R^{S\&P500}_{t+1})}}. \tag{20}$$

The evidence in Figure 5 and in Table VIII suggests once again that the misclassification induces, on average, to the overestimation of the parameter of interest.[13]

[13] The expression above for $\hat{\gamma}$ throws some light on the reasons for the upward bias: the random misclassification lowers the denominator, as the consumption growth series of the non-shareholders I randomly throw in is less correlated with asset returns and exhibit a lower variance than the corresponding series of shareholders.

Table VIII. Monte Carlo simulation of the properties of risk aversion estimators with misclassified shareholders and non-shareholders

5,000 Monte Carlo iterations for various values of the shocks to consumption growth ($\sigma_{\Delta ln(\varepsilon^{s}_{h,t+1})}$ and $\sigma_{\Delta ln(\varepsilon^{ns}_{h,t+1})}$) and of the rate of replacement of shareholders with non-shareholders (noise). The risk aversion estimators considered are the one that minimizes the distance of the IMRS from the Hansen and Jagannathan bound, the one based on the calibration of the non-linear C-CAPM, and the one based on the calibration of the unconditional log-linearized C-CAPM.

		Noisily predicted shareholders (γ^{nss})	Predicted shareholders (γ^{s})	($\gamma^{nss} - \gamma^{s}$)
$\sigma_{\Delta ln\varepsilon^{s}} = 0.505$; $\sigma_{\Delta ln\varepsilon^{ns}} = 0.479$; noise: 35%				
Hansen and	Mean	6.39	4.25	2.14
Jagannathan bound	St. dev.	(0.79)	(0.16)	(0.78)
	Median	6.15	4.25	1.90
Calibration of the	Mean	21.23	16.36	4.88
non-linear model	St. dev.	(10.66)	(8.32)	(6.17)
	Median	19.50	15.00	4.00
Calibration of the	Mean	57.45	46.36	11.10
log-linear model	St. dev.	(68.21)	(66.29)	(51.70)
	Median	45.51	37.22	6.81
$\sigma_{\Delta ln\varepsilon^{s}} = 0.505$; $\sigma_{\Delta ln\varepsilon^{ns}} = 0.479$; noise: 45%				
Hansen and	Mean	7.52	4.58	2.66
Jagannathan bound	St. dev.	(0.66)	(0.19)	(0.64)
	Median	7.04	4.85	2.55
Calibration of the	Mean	22.31	16.16	6.15
non-linear model	St. dev.	(11.52)	(8.51)	(7.23)
	Median	20.00	14.50	5.00
Calibration of the	Mean	59.52	45.87	13.65
log-linear model	St. dev.	(71.71)	(63.23)	(56.54)
	Median	47.28	36.09	8.53
$\sigma_{\Delta ln\varepsilon^{s}} = 0.505$; $\sigma_{\Delta ln\varepsilon^{ns}} = 0.479$; noise: 25%				
Hansen and	Mean	5.78	4.62	1.16
Jagannathan bound	St. dev.	(0.31)	(0.18)	(0.26)
	Median	5.75	4.65	1.15
Calibration of the	Mean	19.13	16.01	3.12
non-linear model	St. dev.	(9.38)	(8.36)	(4.37)
	Median	17.50	14.50	2.50
Calibration of the	Mean	63.61	49.14	14.47
log-linear model	St. dev.	(68.21)	(66.29)	(51.70)
	Median	45.51	37.22	6.81

Table VIII. Continued

		Noisily predicted shareholders (γ^{nss})	Predicted shareholders (γ^{s})	($\gamma^{nss} - \gamma^{s}$)
$\sigma_{\Delta ln\varepsilon^s} = 0.235$; $\sigma_{\Delta ln\varepsilon^{ns}} = 0.240$; noise: 35%				
Hansen and	Mean	7.00	4.29	2.71
Jagannathan bound	St. dev.	(0.64)	(0.16)	(0.62)
	Median	6.95	4.30	2.65
Calibration of the	Mean	19.13	16.01	3.12
non-linear model	St. dev.	(9.94)	(8.36)	(4.37)
	Median	17.50	14.50	2.50
Calibration of the	Mean	54.16	47.36	6.80
log-linear model	St. dev.	(48.69)	(48.23)	(23.79)
	Median	44.89	37.00	6.81
$\sigma_{\Delta ln\varepsilon^s} = 0.235$; $\sigma_{\Delta ln\varepsilon^{ns}} = 0.240$; noise: 35%				
Hansen and	Mean	7.19	4.92	2.27
Jagannathan bound	St. dev.	(0.48)	(0.18)	(0.43)
	Median	7.15	4.90	2.20
Calibration of the	Mean	21.44	16.27	5.17
non-linear model	St. dev.	(10.68)	(8.40)	(5.53)
	Median	20.00	15.00	4.00
Calibration of the	Mean	58.73	49.35	9.39
log-linear model	St. dev.	(60.13)	(56.43)	(40.93)
	Median	45.83	37.31	6.68
$\sigma_{\Delta ln\varepsilon^s} = 0.235$; $\sigma_{\Delta ln\varepsilon^{ns}} = 0.240$; noise: 35%				
Hansen and	Mean	8.22	3.40	2.82
Jagannathan bound	St. dev.	(0.38)	(0.24)	(0.26)
	Median	8.20	5.40	2.80
Calibration of the	Mean	22.37	16.24	6.13
non-linear model	St. dev.	(11.34)	(8.48)	(4.96)
	Median	20.50	15.00	5.00
Calibration of the	Mean	53.09	47.93	5.16
log-linear model	St. dev.	(44.24)	(46.19)	(21.85)
	Median	45.57	37.54	7.03

Overall, it seems fair to conclude that the misclassification may be a source of upward bias. Hence, the true risk aversion coefficient can be expected to be even lower than the estimated values. Some intuition for the upward bias can be obtained by noticing that including the consumption of some randomly-chosen non-shareholders in the measure of the IMRS and excluding that of some randomly-chosen shareholders may reduce the volatility of the stochastic discount factor and its correlation with stock returns. This implies that, if anything, the misclassification plays against finding differences between shareholders and the representative agent.

References

Arrow, K. (1971) *Essays in the Theory of Risk-Bearing*, North Holland, Amsterdam.

Attanasio, O., Banks, J., and Tanner, S. (2002) Asset holding and consumption volatility, *Journal of Political Economy* **110**, 771–792.

Attanasio, O. and Weber, G. (1995) Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey, *Journal of Political Economy* **103**, 1121–1157.

Attanasio, O. and Weber, G. (1993) Consumption growth, the interest rate and aggregation, *Review of Economic Studies* **60**, 631–649.

Banks, J., Blundell, R., and Brugiavini, A. (1999) Risk pooling, precautionary saving and consumption growth, Working Paper No. 99/19, Institute for Fiscal Studies.

Bertaut, C. and Starr-McCluer, M. (2002) Household portfolios in the United States, in L. Guiso, M. Haliassos and T. Japelli eds.: *Household Portfolios*, The MIT Press, Cambridge, Massachusetts.

Breeden, D. (1979) An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* **7**, 265–296.

Brock, W. (1982) Asset prices in a production economy, in J. J. McCall ed.: *The Economics of Information and Uncertainty*, University of Chicago Press, Chicago, Illinois.

Campbell, J. and Cochrane, J. (1995) By force of habit: A consumption-based explanation of aggregate stock market behavior, NBER Working paper No. 4995.

Constantinides, G. (1982) Intertemporal asset pricing with heterogeneous consumers and without demand aggregation, *Journal of Business* **55**, 253–267.

Constantinides, G. (1990) Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* **98**, 519–543.

Epstein, L. and Zin, S. (1989) Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* **57**, 937–969.

Epstein, L. an Zin, S. (1991) Substitution, risk aversion and the temporal behavior of consumption and asset returns: An empirical analysis, *Journal of Political Economy* **99**, 263–286.

Fama, E. and French K. (1989) Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* **47**, 511–552.

Friend, I. and Blume M. (1975) The demand for risky asset, *American Economic Review* **65**, 900–922.

Gallant, R., Hansen, L., and Tauchen, G. (1990) Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution, *Journal of Econometrics* **45**, 141–179.

Haliassos, M. and Bertaut, C. (1995) Why do so few hold stocks? *Economic Journal* **105**, 1110–1129.

Hall, R. (1988) Intertemporal substitution in consumption, *Journal of Political Economy* **96**, 339–357.

Hansen, L. and Jagannathan, R. (1991) Implications of security market data for models of dynamic economies., *Journal of Political Economy* **99**, 225–262.

Hansen, L. and Singleton, K. (1982) Generalized instrumental variables estimation of non-linear rational expectation models, *Econometrica* **50**, 1269–1286.

Hansen, L. and Singleton, K. (1983) Stochastic consumption, risk aversion, and the temporal behavior of asset returns, *Journal of Political Economy* **91**, 249–266.

Hansen, L. and Singleton, K. (1984) Errata (to the 1982 paper) *Econometrica* **52**, 267–268.

Heaton, J. (1995) An empirical investigation of asset pricing with temporally dependent preference specification, *Econometrica* **63**, 681–717.

Heaton, J. and Lucas, D. (1996) Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* **104**, 443–487.

Kydland, F. and Prescott, E. (1982) Time to build and aggregate fluctuations, *Econometrica* **50**, 1345–1370.

Kocherlakota, N. (1996) The equity premium: It is still a puzzle, *Journal of Economic Literature* **24**, 42–71.

Lintner, J. (1965) The valuation of risk assets and the selection of risky investment in stock portfolios and capital budgets, *Review of Economics and Statistics* **47**, 13–37.

Lucas, R. (1978) Asset prices in an exchange economy, *Econometrica* **46**, 1429–1445.

Mankiw, G. and Zeldes, S. (1991) The consumption of stockholders and non-stockholders, *Journal of Financial Economics* **29**, 97–112.

Mehra, R. and Prescott, E. (1985) The equity premium. A puzzle, *Journal of Monetary Economics* **15**, 145–161.

Merton, R. (1973) An intertemporal asset pricing model, *Econometrica* **41**, 867–887.

Newey, W. and West, K. (1987) A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* **55**, 703–708.

Poterba, J. and Samwick, A. (1995) Stock ownership patterns, stock market fluctuations, and consumption, *Brookings Papers on Economic Activity* **2**, 295–372.

Sabelhaus, J. (1998) Using the consumer expenditure survey for research on saving and consumption, unpublished working paper, Congressional Budget Office.

Sharpe, W. (1964) Capital asset prices: A theory of capital market equilibrium under conditions of risk, *Journal of Finance* **19**, 425–442.

Vissing-Jørgensen, A. (1999) Limited stock market participation, unpublished working paper, Northwestern University.

Vissing-Jørgensen, A. (2002) Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* **110**, 825–853.

Copyright of Review of Finance is the property of Kluwer Academic Publishing / Business 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.