Limited Asset Market Participation and the Elasticity of Intertemporal Substitution

Annette Vissing-Jørgensen

University of Chicago, National Bureau of Economic Research, and Centre for Economic Policy Research

The paper presents empirical evidence based on the U.S. Consumer Expenditure Survey that accounting for limited asset market participation is important for estimating the elasticity of intertemporal substitution. Differences in estimates of the EIS between asset holders and non–asset holders are large and statistically significant. This is the case whether estimating the EIS on the basis of the Euler equation for stock index returns or the Euler equation for Treasury bills, in each case distinguishing between asset holders and non–asset holders as best as possible. Estimates of the EIS are around 0.3–0.4 for stockholders and around 0.8–1 for bondholders and are larger for households with larger asset holdings within these two groups.

I. Introduction

The elasticity of intertemporal substitution (EIS) is one of the central determinants of households' intertemporal consumption choices. In the certainty case it measures the (inverse of) the elasticity of the marginal rate of substitution between consumption at two different dates with

This paper is a revised version of the first chapter of my Massachusetts Institute of Technology thesis. I thank my advisors Daron Acemoglu, Olivier Blanchard, and especially Ricardo Caballero for their insights and advice. Seminar and conference participants and discussants at MIT, Princeton, Rochester, Chicago, Northwestern, Kellogg, Stanford, New York University, Columbia, Wisconsin, University College London, Tilburg, Toulouse, Tel Aviv, Milan, the NBER Summer Institute, and the American Economic Association meetings provided very useful comments and suggestions. Thanks also to the *Journal* editors and referees, whose suggestions substantially improved the paper. All errors are mine. Financial support from the Danish Research Academy, the Sloan Foundation, and the University of Chicago is gratefully acknowledged.

[*Journal of Political Economy,* 2002, vol. 110, no. 4]

respect to the ratio of the consumption values. It can be computed as minus the elasticity of the ratio of consumption in period t to consumption in period $t+1$ with respect to the relative price of consumption in the two periods:

$$\text{EIS} = -\frac{d(C_t/C_{t+1})}{d(1+R_{f,t})}\frac{1+R_{f,t}}{C_t/C_{t+1}} = \frac{d[\ln (C_{t+1}/C_t)]}{d[\ln (1+R_{f,t})]}, \tag{1}$$

where $R_{f,t}$ is the real net return on the riskless asset in period t. Under uncertainty, one can define utility over consumption in this period and the certainty equivalent of future consumption following Epstein and Zin (1989). The EIS can then be defined on the basis of the ratio of current-period consumption to the certainty equivalent of future consumption. Under the assumption that asset returns and next-period consumption are joint lognormally distributed conditional on information known in this period, Hansen and Singleton (1983) show for the case of constant relative risk aversion (CRRA) utility that the EIS can be calculated as

$$\text{EIS} = \frac{dE_t[\ln (C_{t+1}/C_t)]}{dE_t[\ln (1+R_{i,t})]}, \tag{2}$$

where $R_{i,t}$ is the real net return in period t on any asset i in which the consumer has an interior position. In other words, the EIS determines how much consumers change their expected consumption growth rate in response to changes in the expected return to any such asset. Attanasio and Weber (1989) give a similar derivation for the Epstein-Zin utility case and conclude that (2) remains valid as an approximation for these more general preferences.

Determining the sensitivity of consumption to interest rates is important for a host of economic issues and policies. For example, (*a*) the response of household savings to changes in after-tax interest rates depends crucially on the EIS, (*b*) the effect of expansionary fiscal policy is weaker if the EIS is large, and (*c*) the ability of real business cycle models to fit the data depends on the EIS (see, e.g., Jones, Manuelli, and Siu [2000] regarding point *c*).

A quite large empirical literature has been devoted to estimating the EIS. Most work has focused on estimating it assuming CRRA preferences and using the log-linearized Euler equation derived by Hansen and Singleton (1983). Hall (1988) summarized the evidence up to the late 1980s by concluding that the EIS is unlikely to be much above 0.1 and may well be zero. While some subsequent studies (e.g., Attanasio and Weber 1993) have found positive and significant values of the EIS, it seems fair to say that there is no consensus on whether it is significantly above zero and, if so, what its value is.

In this paper I use micro data from the U.S. Consumer Expenditure Survey (CEX) to argue that accounting for limited asset market participation is crucial for obtaining consistent estimates of the EIS. The Euler equation should hold for a given household only if the household holds a nonzero position (positive or negative) in the asset. There is no theoretical reason for expecting households that do not hold a given asset to adjust their consumption growth rate in response to predictable changes in the return on that asset. In that case, including the consumption of non–asset holders in Euler equation estimations will lead to inconsistent estimates of the EIS. In particular, if the consumption growth of non–asset holders does not comove with predictable asset return changes at all, then EIS estimates based on the consumption of all households will be substantially downward-biased.[1] To test this prediction, I estimate log-linearized Euler equations involving Treasury bill returns as well as log-linearized Euler equations involving the return on the New York Stock Exchange (NYSE) stock market index. In each case I distinguish, as best as possible given the CEX data, between households that hold the asset in question and households that do not. I find estimates of the EIS around 0.3–0.4 for stockholders and around 0.8–1 for bondholders, with larger values for the third of households within these two groups that have the largest holdings of the asset. For nonstockholders and nonbondholders, EIS estimates are small and insignificantly different from zero. Interestingly, in their paper in this issue, Attanasio, Banks, and Tanner (2002) also find support for the limited asset market participation theory using data for the United Kingdom. These findings are consistent with earlier work by Attanasio and Browning (1995) showing that the EIS is increasing in consumption, given the large correlation of consumption and wealth and the fact that wealth is known to be a strong predictor of participation in the stock and bond markets.

It is important to emphasize that the differences in EIS estimates between asset holders and non–asset holders should not be interpreted as evidence of heterogeneity in the EIS across households. Since the Euler equation for a given asset return cannot be expected to hold for households that do not hold a position in that asset, EIS estimates for non–asset holders are not consistent estimates of the EIS for these households. They are shown only to reconcile my findings of fairly large values

[1] For example, suppose that households have CRRA preferences and consumption growth rates and asset returns are conditionally joint lognormal. If the consumption growth of non–asset holders does not covary with the instruments used to predict the asset return, it can be shown that EIS estimates based on estimation of conditional Euler equations using the consumption growth rates of all households are only factor λ of the true value, where λ is the fraction of asset holders in the population. The derivation of this is based on cross-sectional aggregation of individual Euler equations as in the empirical section below.

of the EIS for asset holders with the findings of earlier studies. In fact it may be that the EIS is similar for asset holders and non–asset holders. In a recent paper, Gross and Souleles (2002) find a significant negative relation between the interest rate on credit cards and the amount of credit card borrowing, suggesting that the EIS is significantly positive for credit card borrowers as well.

The idea that limited participation in asset markets matters for consumption and asset returns was originally proposed by Mankiw and Zeldes (1991). They estimated unconditional Euler equations for stockholders and nonstockholders using data from the Panel Study of Income Dynamics (PSID). On the basis of unconditional Euler equations, they found large differences in relative risk aversion estimates between the two groups, although the estimate remained as high as 35 for the richest group of stockholders. Brav and Geczy (1996) and the extended version by Brav, Constantinides, and Geczy (2002; this issue) confirm Mankiw and Zeldes's findings for unconditional Euler equations using CEX data. They show that risk aversion estimates decline as they look at still wealthier layers of asset holders. In the present paper, I focus on the effects of limited participation for estimates of the elasticity of intertemporal substitution rather than estimates of risk aversion. Mankiw and Zeldes and Brav et al. do not provide standard errors for their estimates. However, bootstrap methods show that the standard errors of risk aversion estimates based on unconditional Euler equations are very large, especially when one uses a relatively short time period as is available from the CEX (see Vissing-Jørgensen 1999). This suggests that adding information about predictable movements in expected consumption growth rates and expected asset returns is valuable for improving the precision of the estimates. If, however, preferences do not have the CRRA form, the coefficient of relative risk aversion does not equal the inverse of the EIS. More specifically, in the case of Epstein-Zin preferences, estimation of conditional log-linearized Euler equations, with log consumption growth regressed on log asset returns, provides estimates of the EIS but is not informative about risk aversion (see, e.g., Attanasio and Weber 1989). Therefore, while the findings of the present paper suggest that accounting for limited asset market participation is crucial for obtaining reasonable values of the EIS, it is too early to precisely determine the extent to which it helps resolve the equity premium puzzle.

II. Empirical Results Based on the Consumer Expenditure Survey

A. *Data*

The available data from the CEX cover the period 1980:1–1996:1. In each quarter, approximately 5,000 households are interviewed. Each household is interviewed five times. The first time is practice, and the results are not in the data files. The interviews are three months apart, and when interviewed, households are asked to report consumption for the previous three months. Financial information is gathered only in the fifth quarter. Aside from attrition, the sample is representative of the U.S. population. Attrition is quite substantial, with only about 70 percent of the households having data for all four interviews.

1. Asset Holder Status

The CEX contains information about four categories of financial assets. Households are asked for their holdings of "stocks, bonds, mutual funds, and other such securities," "U.S. savings bonds," "savings accounts," and "checking accounts, brokerage accounts, and other similar accounts."

I perform estimations first based on the return on the NYSE stock market index and then based on the return on Treasury bills. Thus I would like to separate households into those that own stocks versus those that do not, as well as into those that hold bonds, Treasury bills, or similar assets versus those that do not. As is clear from the categories above, a perfect separation of households is not possible. In general, inability to perfectly identify asset holders from non–asset holders biases against finding differences in EIS estimates for the two groups. I shall refer to households with positive responses to the category stock, bonds, mutual funds, and other such securities as stockholders and focus on this distinction when estimating Euler equations for the stock index return. As for bondholdings, I shall refer to households with positive responses to the stock, bonds, and mutual funds category or to the U.S. savings bond category as bondholders (thus my sets of stockholders and bondholders differ by the group of households that own U.S. savings bonds). Again, not all of those with positive holdings of stocks, bonds, or mutual funds necessarily hold bonds or bond mutual funds, but many likely do. I do not include households with positive holdings in savings accounts in the bondholder category but return to these households after the main estimation results.

The Euler equation involving consumption in periods t and $t+1$ should hold for those that hold the asset at the beginning of period t. Therefore, asset-holding status must be defined on the basis of stockholdings at the beginning of period t. Two additional CEX variables are used for this purpose. The first one reports whether the household holds the same amount, more, or less of the asset category compared to a year ago. The second one reports the dollar difference in the estimated market value of the asset category held by the household last month compared with the value of the asset category held a year ago last month. I define a household as holding an asset category at the beginning of period t if it (1) reports holding the same amount of the asset as a year ago and holds a positive amount at the time of the interview (the fifth interview), (2) reports having lower holdings of the asset than a year ago, or (3) reports having had an increase in its holdings of the asset but by a dollar amount less than the reported holdings at the time of the question.[2] On this basis I classify 21.75 percent of households as stockholders and 31.40 percent as bondholders.

It is known from, for example, the Survey of Consumer Finances that many households hold stocks or bonds only in their pension plan. Whether they should be considered stockholders/bondholders or not depends on the type of pension plan. In a defined contribution pension plan, households can (within some limits) adjust their contributions and allocations and thus ensure that the Euler equation for the asset return is satisfied. Unfortunately, it is not possible to determine whether households with defined contribution plans report their stockholdings and bondholdings in these plans when answering the CEX questions. The percentage of stockholders in the CEX is smaller than in other sources.[3] This may indicate that some households with stockholdings in defined contribution plans do not report them. This will lead them to be miscategorized as nonstockholders, a problem that may also occur for bondholdings. Again, this should bias against finding differences between asset holders and non–asset holders.

In addition to the split between stockholders and nonstockholders and the split between bondholders and nonbondholders, the set of stockholders and the set of bondholders are split into three layers of approximately equal size on the basis of dollar amounts reported. Consistent with the definition of asset-holding status, this was done on the basis of initial holdings. These were calculated as current holdings minus the change in holdings during the current period. Transactions costs may be less important for wealthier households, making it more likely that their Euler equations for stocks and bonds hold. Furthermore, wealthier households typically hold better-diversified equity portfolios. The return on their equity portfolio is thus closer to the stock market index return used here. The bottom layer of stockholders consists of those reporting initial holdings of \$2–\$3,587 in real 1982–84 dollars,

[2] Around 600 households in my final sample of 34,310 households report an increase in their holdings of stocks, bonds, and mutual funds but do not report their current holdings. Most of these households are likely to have held these assets a year ago, and I therefore put them in the stockholder category. Similarly, around 400 households report an increase in their holdings of U.S. savings bonds but do not report their current holdings. I classify these as bondholders. A small number of households report an increase in their holdings of stocks, bonds, and mutual funds or of U.S. savings bonds larger than the value of the reported end-of-period holdings. I classify these as non–asset holders.

[3] This is the case whether the CEX weights are used or not. Consistent with data from other sources, the proportion of stockholders is upward-trending during the sample.

with the consumer price index for total consumption of urban households used to deflate the nominal values. The middle and top layers are those with real initial holdings of \$3,587–\$20,264 and above \$20,264, respectively. For interviews conducted from 1991 onward, about 5 percent of households report holdings of stocks, bonds, and mutual funds of \$1. I contacted the Bureau of Labor Statistics to determine whether this was a coding error, but it was not sure how to interpret the \$1 answers. Since all the households reporting \$1 asset holdings answer the question comparing current holdings to holdings a year ago, it is likely that they are holding such assets. I therefore include them as stockholders when doing the stockholder-nonstockholder classification (assuming they satisfy the criteria outlined above). However, since the \$1 households cannot be classified by layer of stockholding, I exclude them in estimations based on layers of stockholders. Thus the total number of stockholders used when defining layers is smaller than the total number of stockholders when not doing this split. As for the three layers of bondholders, they were calculated on the basis of the sum of the holdings of stocks, bonds, and mutual funds and U.S. savings bonds, again excluding households reporting \$1 holdings of the first category. The least wealthy third had holdings below \$967, the middle third had holdings between \$967 and \$9,881, and the top third held more than \$9,881 combined in these asset categories.

2. Consumption Measure and Sample Choice

The consumption measure used is nondurables and some services aggregated as carefully as possible from the disaggregate CEX consumption categories to match the definitions of nondurables and services in the National Income and Product Accounts (NIPA). The service categories excluded are housing expenses (but not costs of household operations), medical care costs, and education costs. This is done since these three types of costs have a substantial durable component. Attanasio and Weber (1995) use a similar definition of consumption. When durables are left out, it is implicitly assumed that utility is separable in durables and nondurables/services. Nominal consumption values are deflated by the Bureau of Labor Statistics deflator for nondurables for urban households.

As discussed in more detail below, I use semiannual consumption growth rates, defined as

$$\frac{C_{m+6} + C_{m+7} + C_{m+8} + C_{m+9} + C_{m+10} + C_{m+11}}{C_m + C_{m+1} + C_{m+2} + C_{m+3} + C_{m+4} + C_{m+5}},$$

where m refers to month m. As in Zeldes (1989), extreme outliers are dropped under the assumption that they reflect reporting or coding errors. Specifically, I drop observations for which the consumption growth ratio stated is less than 0.2 or above five (35 observations). Results are similar if I keep these observations. In addition, nonurban households (missing for part of the sample) and households residing in student housing are dropped, as are households with incomplete income responses. Furthermore, I drop households that report a change in age of household head between any two interviews different from zero or one year. These exclusions are standard. More drastically, I drop all consumption observations for households interviewed in 1980 and 1981 since the quality of the CEX consumption data is lower for this period. For example, figure A.1 in Parker (2000) shows that the ratio of food consumption to income in the CEX is substantially higher before 1982 and that a similar pattern is not present in the PSID or the NIPA. Attanasio and Weber (1995) show that the share of food in nondurable consumption was much higher in 1980 and 1981 than subsequently.

Finally, because the financial information is reported in interview 5 and because I wish to calculate consumption growth values by household, households must be matched across quarters. Therefore, I drop households for which any of interviews 2–5 are missing. Matching households across interviews creates problems around the beginning of 1986 since sample design and household identification numbers were changed, with no records being kept of which new household identification numbers correspond to which old ones. I therefore exclude households that did not finish their interviews before the identification number change. This implies that no observations are available for seven months around the identification change. I take this into account when programming corrections for autocorrelation. The final sample consists of 34,310 semiannual consumption growth observations.

Table 1 gives the average number of observations per month for the various household groups, along with the mean and standard deviation of their semiannual average log consumption growth (not annualized). The average number of observations for stockholders per month is 47, and the average number of observations for nonstockholders is 170. There are, on average, 68 bondholders and 149 nonbondholders per month. Data are also shown for a restricted sample consisting of only single-individual households. This sample has substantially fewer observations. Therefore, it is not possible to consider layers of stockholders and bondholders in this case. The summary statistics on consumption in table 1 show that log consumption growth is higher for stockholders and bondholders than for nonstockholders and nonbondholders.

TABLE 1
SUMMARY STATISTICS, CEX DATA, 1982–96

Group	Mean Number of Observations per Month	Mean of $(1/H_t)\sum_{h=1}^{H_t}\Delta \ln C_{t+1}^h$ over the Sample	Standard Deviation of $(1/H_t)\sum_{h=1}^{H_t}\Delta \ln C_{t+1}^h$ over the Sample
	A. Semiannual Data, All Household Sizes		
All	217	.002	.024
Stockholders	47	.014	.041
Nonstockholders	170	−.001	.028
Bottom stockholder layer	11	.022	.088
Middle stockholder layer	11	.002	.101
Top stockholder layer	11	.014	.102
Bondholders	68	.010	.034
Nonbondholders	149	−.001	.030
Bottom bondholder layer	17	.005	.069
Middle bondholder layer	17	.012	.072
Top bondholder layer	17	.008	.073
	B. Semiannual Data, Single-Individual Households		
All	51	−.001	.045
Stockholders	10	.009	.098
Nonstockholders	41	−.003	.052
Bondholders	13	.002	.085
Nonbondholders	39	−.002	.054

NOTE.—The variable $(1/H_t)\sum_{h=1}^{H_t}\Delta \ln C_{t+1}$ is seasonally adjusted using seasonal dummies. The seasonally adjusted value is the sample mean of the series plus the residual from a regression on 12 dummies.

3. Asset Returns and Other Data Used

Monthly NYSE value-weighted returns are used as the stock return measure and monthly Treasury bill returns as the measure of nominally riskless returns. The consumer price index for total consumption of urban households is used to calculate real returns. Semiannual returns are aggregated up from the real monthly returns. As instruments for the log stock return and the log Treasury bill return I use the log dividend-price ratio, the lagged log real value-weighted NYSE return, the lagged log real Treasury bill return, the lagged government bond horizon premium, and the lagged corporate bond default premium. The choice of instruments is discussed further below. The dividend-price ratio, the bond horizon premium, and the bond default premium are based on data from Ibbotson Associates (1997). The dividend-price ratio used is the ratio of dividends over the previous 12 months to the current price (the Standard & Poor's 500 index). The bond horizon premium is defined as

$$\frac{1 + R_t^{\text{long-term government bonds}}}{1 + R_t^{\text{short-term government bonds}}},$$

where "long-term" means 20 years to maturity and "short-term" means approximately one month to maturity. The bond default premium is defined as

$$\frac{1 + R_t^{\text{long-term corporate bonds}}}{1 + R_t^{\text{long-term government bonds}}},$$

where long-term again means 20 years to maturity. The monthly values for the bond horizon and default premia are aggregated multiplicatively to semiannual values.

B. Econometric Issues

1. Timing

The fact that households are interviewed every three months for a year and in each interview report consumption for the previous quarter leaves open a choice of data frequency for defining consumption growth rates. I use semiannual consumption growth rates as defined above and thus have one consumption growth observation per household interviewed. The first motivation for using semiannual rather than quarterly consumption growth rates is that households may not reoptimize and optimally adjust their consumption every quarter (or, even if they do, may not do it close to the interview dates). In the context of risk aversion estimation, Lynch (1996), Daniel and Marshall (1997), and Gabaix and Laibson (2002) discuss how this leads to a downward-biased prediction for the equity premium and thus by implication an upward-biased risk aversion estimate, when individual consumption is aggregated across households. Using consumption observations further apart, on average six months for semiannual consumption growth rates, alleviates the problem. A second argument in favor of longer time horizons is measurement error. If measurement errors have an additive component, this will tend to cancel over longer periods. Furthermore, measurement errors may be negatively correlated with the true consumption value. Daniel and Marshall suggest that measurement errors that are negatively correlated with the true value can occur if consumption innovations are only gradually incorporated into reported consumption over time. This will bias the covariance of consumption growth rates and asset returns downward, but less so if the consumption observations are farther apart. Consistent with the concerns given, in my exercise, results based on quarterly data were much weaker than those based on semiannual data. Results based on the change from first interview quarterly consumption to last interview quarterly consumption were quite similar to those based on semiannual consumption growth rates.

Using semiannual consumption growth rates raises the issue of precisely which asset return to use. For semiannual data, is the relevant asset return $(1 + R_m)(1 + R_{m+1}) \cdots (1 + R_{m+5})$ or $(1 + R_{m+6})(1 + R_{m+7}) \cdots (1 + R_{m+11})$ or something in between? Suppose that consumers sell their assets (stocks or bonds) at the beginning of the month in which they would like to consume. Then shifting consumption from period m to $m + 6$ would imply a gross return of $(1 + R_m)(1 + R_{m+1}) \cdots (1 + R_{m+5})$, whereas shifting consumption from period $m + 5$ to $m + 11$ would yield a return of $(1 + R_{m+5})(1 + R_{m+6}) \cdots (1 + R_{m+10})$. Thus the relevant asset return would be a weighted average of all of $(1 + R_m)$, $(1 + R_{m+1})$, $\ldots$, $(1 + R_{m+10})$. For simplicity I use the middle six months of relevant interest rates $(1 + R_{m+2})(1 + R_{m+3}) \cdots (1 + R_{m+7})$. Results were similar when I used $(1 + R_{m+3})(1 + R_{m+4}) \cdots (1 + R_{m+8})$.

While each household is interviewed three months apart, the interviews are spread out over the quarter, implying that there will be households interviewed in each month of the sample. Thus the data frequency is monthly. This implies that observations of consumption growth for adjacent months will involve partially overlapping time periods and thus partially overlapping expectational errors. As a consequence, the error term in the log-linearized model will have an MA(5) component when semiannual observations of consumption growth are used.

Autocorrelation raises the question of which lags of interest rates and other variables are valid instruments. The error term has an expectational error component and a measurement error component. Since asset returns and other aggregate financial variables are likely to be uncorrelated with the measurement error component of the error term, autocorrelation due to measurement error does not invalidate lags of asset returns or other aggregate financial variables as instruments. The autocorrelation in the error term due to overlapping expectational errors implies that asset return lags six and farther back are valid instruments. Since R_m and R_{m+1} may be partially relevant as right-hand-side variables as discussed above, the asset return instruments used are lagged farther so that there is no overlap. Thus $(1 + R_{m-6})(1 + R_{m-5}) \cdots (1 + R_{m-1})$ is used as the instrument. As for the specific asset returns used, I shall return to when stock returns and Treasury bill returns are used when discussing the results. Similar considerations lead to using the dividend-price ratio at the beginning of period m as an instrument, as well as the bond horizon premium and bond default premium over the period $m - 6$ to $m - 1$.

2. Measurement Error

The conditions on measurement error under which consistent estimates of the EIS can be obtained on the basis of Euler equation estimation are quite strict. A sufficient condition is that the measurement error in individual consumption is multiplicative and independent of the true consumption level and of asset returns as well as any instruments used in the estimation. This is the case whether the Euler equations are log-linearized or not. For simplicity, the argument below therefore uses the nonlinear Euler equations.

To be specific, suppose first that we have a long time series of consumption observations for a household h and wish to test the consumption capital asset pricing model. Let $C_t^{h,*}$ be the true consumption of household h in period t, and assume that observed consumption is given by $C_t^h = C_t^{h,*}\epsilon_t^h$, where ϵ_t^h is the measurement error. Under CRRA utility with risk aversion parameter γ (= 1/EIS), the true Euler equation for household h and asset i is

$$E_t\left[\beta\left(\frac{C_{t+1}^{h,*}}{C_t^{h,*}}\right)^{-\gamma}(1+R_{i,t})\right] = 1. \tag{3}$$

However, our estimates $\hat{\beta}$ and $\hat{\gamma}$ will be based on sample equivalents of

$$E_t\left[\beta\left(\frac{C_{t+1}^{h}}{C_t^{h}}\right)^{-\gamma}(1+R_{i,t})\right] = 1 \Leftrightarrow E_t\left[\beta\left(\frac{C_{t+1}^{h,*}\epsilon_{t+1}^h}{C_t^{h,*}\epsilon_t^h}\right)^{-\gamma}(1+R_{i,t})\right] = 1. \tag{4}$$

As an example, in the case of generalized method of moments (GMM) estimation with two asset returns and two instruments, one of which is a column of ones, the estimates $\hat{\beta}$ and $\hat{\gamma}$ satisfy the sample equivalents of (the unconditional version of) equation (4) above with equality. If ϵ_{t+1}^h and ϵ_t^h are independent of $C_{t+1}^{h,*}$, $C_t^{h,*}$, and $R_{i,t}$, (4) implies

$$E\left[\left(\frac{\epsilon_{t+1}^h}{\epsilon_t^h}\right)^{-\gamma}\right]E\left[\beta\left(\frac{C_{t+1}^{h,*}}{C_t^{h,*}}\right)^{-\gamma}(1+R_{i,t})\right] = 1. \tag{5}$$

It follows that the estimator of β will be inconsistent by the factor $E[(\epsilon_{t+1}^h/\epsilon_t^h)^{-\gamma}]$, whereas γ will be consistently estimated. The same conclusion would be reached by considering the other equation used in the GMM estimation, assuming independence of measurement errors and the instrument. If measurement errors are independently and identically distributed lognormal, $\ln \epsilon_t^h \sim N(\mu_\epsilon, \sigma_\epsilon^2)$, $\hat{\beta}$ will thus be inconsistent by the factor $E[(\epsilon_{t+1}^h/\epsilon_t^h)^{-\gamma}] = \exp(\gamma^2\sigma_\epsilon^2)$.[4]

[4] Similarly, $\widehat{\ln \beta}$ in the log-linearized model will be inconsistent by the quantity $\gamma^2\sigma_\epsilon^2$. However, since the intercept in the log-linearized model includes higher-order moments, it is not possible to accurately estimate β from this intercept even in the absence of measurement error.

If we do not have a long time series of consumption for each agent and instead average marginal rates of substitution over consumers within each period, it can be shown that as the number of households in the cross section at each date goes to infinity, γ is still consistently estimated and the inconsistency in β is as given above. To see this, suppose that at least two observations are available for each household such that C^h_{t+1}/C^h_t can be calculated. When the marginal rates of substitution are aggregated cross-sectionally, with a time series of cross sections of C^h_{t+1}/C^h_t observations, $\hat{\beta}$ and $\hat{\gamma}$ will be based on the sample equivalents of

$$E_t\left[\beta\frac{1}{H}\sum_h\left[\left(\frac{C^h_{t+1}}{C^h_t}\right)^{-\gamma}\right](1+R_{i,t})\right]=1$$
$$\Leftrightarrow E_t\left[\beta\frac{1}{H}\sum_h\left[\left(\frac{C^{h,*}_{t+1}\epsilon^h_{t+1}}{C^{h,*}_t\epsilon^h_t}\right)^{-\gamma}\right](1+R_{i,t})\right]=1 \tag{6}$$

or, as $H\rightarrow\infty$,

$$E_t\left[\beta E_h\left[\left(\frac{C^h_{t+1}}{C^h_t}\right)^{-\gamma}\right](1+R_{i,t})\right]=1$$
$$\Leftrightarrow E_t\left[\beta E_h\left[\left(\frac{C^{h,*}_{t+1}\epsilon^h_{t+1}}{C^{h,*}_t\epsilon^h_t}\right)^{-\gamma}\right](1+R_{i,t})\right]=1. \tag{7}$$

With independence between true consumption and measurement errors, this implies that

$$E_t\left[\beta E_h\left[\left(\frac{C^{h,*}_{t+1}}{C^{h,*}_t}\right)^{-\gamma}\right]E_h\left[\left(\frac{\epsilon^h_{t+1}}{\epsilon^h_t}\right)^{-\gamma}\right](1+R_{i,t})\right]=1. \tag{8}$$

Thus as $H\rightarrow\infty$ (and $T\rightarrow\infty$), $\hat{\gamma}$ will again be consistent, and with independently and identically distributed lognormal measurement errors (in the cross section), $\hat{\beta}$ will be inconsistent by the same factor as above. Since this argument for consistency of $\hat{\gamma}$ even in the presence of measurement error relies on $H\rightarrow\infty$ and the number of observations in the cross section is small in some of the cases considered, I shall also consider a "median household approach," which is likely to be more robust to small numbers of households in the cross sections.[5]

[5] An alternative way to see that measurement error does not prevent consistent estimation of the EIS in either the CRRA or the Epstein-Zin case is to consider the log-linearized Euler equation stated below in which log consumption growth is a left-hand-side variable and lagged consumption growth rates are not used as instruments.

3. Family Size and Seasonality Controls

Following a series of papers in the consumption literature, I assume that family size enters the utility function multiplicatively and thus include $\Delta \ln$ (family size) in the log-linearized Euler equations. The variable $\Delta \ln$ (family size) is defined as the log average family size in the third and fourth interviews minus the log average family size in the first and second interviews. It is, however, not clear that such a simple correction accurately captures family size effects. The literature on equivalence scales considers this issue in detail. Here, I choose to repeat the estimations using households consisting of only a single individual in both periods to see if this affects the results. Single-individual households may also face a much simpler optimization problem, making it easier to detect the relation between consumption growth and stock or Treasury bill returns in the micro data.

I assume that seasonality also enters as a multiplicative factor in the utility function such that seasonal adjustment by dummies is valid in the log-linearized model.

4. Estimation Method

As discussed in the Introduction, I estimate log-linearized Euler equations using standard estimation techniques. To avoid the complicated semiannual notation, use t to denote six-month periods and refer to the subsections above for the precise definitions of semiannual consumption growth rates and family size changes as well as the timing of the asset returns. As earlier, I use h to denote a given household.

Since I have only one consumption growth observation per household, I use the simple cohort technique, discussed in subsection 3 on measurement error, in which the consumption growth observation for a given period is the cross-sectional average of the consumption growth observations for households of a given type in the sample for that period. Notice that as long as one consumption growth observation is available per household, it is still possible to avoid aggregation problems since one can average $\Delta \ln C_{t+1}^h = \ln (C_{t+1}^h / C_t^h)$ across households as opposed to using the log of the ratio of the household consumption averages.[6]

[6] See Brav et al. (2002; this issue) for a detailed analysis of the importance of aggregation issues in the CEX data set.

The log-linearized conditional Euler equations are then, with stockholders' Euler equation for the stock return as an example,

$$\frac{1}{H_t^s}\sum_{h=1}^{H_t^s}\Delta \ln C_{t+1}^{h,s} = \sigma^s \ln(1+R_{s,t}) + \delta_1^s D_1 + \delta_2^s D_2 + \cdots + \delta_{12}^s D_{12} + \alpha^s \frac{1}{H_t^s}\sum_{h=1}^{H_t^s}\Delta \ln(\text{family size})_{t+1}^{h,s} + u_{t+1}^s, \qquad (9)$$

where σ denotes the EIS, $D_1, \ldots, D_{12}$ are seasonal dummies, R_s denotes the real net stock return, and H_t^s denotes the number of stockholders in the cross section at date t.

This equation is valid under CRRA preferences as well as under Epstein-Zin utility. In the CRRA case, δ_j^s is a function of β and of the conditional variances and covariances of the log gross stock return and log consumption growth. In the Epstein-Zin case, δ_j^s includes additional terms involving the variance of and conditional covariances with the return on the total portfolio of assets held by the agents. The error term u_{t+1}^s includes the expectational errors for log consumption growth and log stock returns and the measurement error in log consumption growth. If the conditional variances and covariances in δ_j^s are not constant, the stochastic components enter the error term. This does not cause problems for the estimation as long as these components are uncorrelated with the instruments used.

The estimation method used is linear GMM estimation or, in other words, linear instrumental variables estimation modified to account for autocorrelated error terms of the MA(5) form as well as for heteroscedasticity of arbitrary form. Heteroscedasticity is likely to be present because of a varying number of observations per quarter. Instrumental variables estimation is used rather than ordinary least squares because of endogeneity of asset returns due to inclusion of the expectational error in the error term. For stockholders (also assumed to be bondholders according to my classification), the Euler equations for the stock and Treasury bill returns are also estimated jointly, again using linear GMM estimation. Joint estimation is used to gain efficiency from exploiting cross-equation correlation in error terms caused by correlated expectational errors. Furthermore, it makes it possible to impose identical values for σ to determine whether this leads to rejection of the model according to overidentification tests, which Hansen and Singleton (1983) showed to be the case in aggregate U.S. data. In the joint estimation, the coefficients on the seasonal dummies are allowed to differ for the stock and the bond equations since this is implied by the model when the two returns have different variances or different covariances with log consumption growth.

C. *Results and Discussion*

The results of the instrumental variables estimations of the log-linearized model in equation (9) are shown in tables 2–4. Each table shows estimates of σ for three sets of estimations corresponding to three different sets of instruments. All instrument sets include 12 seasonal dummies and $\Delta \ln$ (family size). In addition, instrument set 1 includes the log dividend-price ratio. Instrument set 2 includes the log dividend-price ratio, the lagged log real stock return, and the lagged log real Treasury bill return. Instrument set 3 includes the log dividend-price ratio, the bond horizon premium, and the bond default premium. See Section II*B*1 for the precise timing of these variables. The dividend-price ratio is well known to be among the best predictors of real stock returns and is also a good predictor of the Treasury bill return for the time period considered here. As for instrument set 2, many previous studies have used lagged asset returns as instruments. For this time period, the contemporaneous dividend-price ratio and the Treasury bill returns are highly correlated, and results based on instrument set 2 are fairly similar if the log dividend-price ratio is excluded. The inclusion of the bond horizon premium and the bond default premium is motivated by the findings of Fama and French (1989) that these premiums have predictive power for stock returns. When I leave out the family size variable, the R^2 from regressing the log real stock return on the variables in the three instrument sets is .117, .143, and .148 for semiannual data. The corresponding R^2 values for the log real Treasury bill return are .673, .674, and .744.

While other variables in the information set should also be uncorrelated with the error terms in the log-linearized Euler equations according to economic theory, the properties of the estimators may deteriorate if weak instruments are included. It is known that the two-stage least squares estimator tends to be biased toward the biased and inconsistent ordinary least squares estimator. The bias becomes progressively worse as the degree of overidentification is increased. This motivates the use of small sets of instruments. Adding in more instruments tended to push the estimates of the EIS closer to zero for all groups.

The results of the estimations are favorable to the limited participation theory. Panel A of table 2 shows the estimation results for the Euler equation for the stock return. For instrument set 1, the EIS is estimated to be 0.299 for stockholders. The estimate is significant at the 5 percent level. The separate estimations for the three layers of stockholders show that the relatively high value of the EIS for stockholders is driven by a higher estimate for the richest layer of stockholders. One interpretation is that wealthier stockholders have a higher elasticity of substitution. Alternatively, it is possible that households with small equity stakes do

TABLE 2

GMM Estimation of Log-Linearized Euler Equations: Real Treasury Bill Return and Real Value-Weighted NYSE Return, Separate Estimations (CEX, 1982–96, Semiannual Data)

	Instrument Set 1		Instrument Set 2			Instrument Set 3		
	$\hat{\sigma}$	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ
	A. Euler Equation for Stocks							
	1. All Household Sizes							
All	.098		.066	.086		.068	.314	
	(.071)		(.062)			(.059)		
Stockholders	**.299**		.281	.260		.200	.018	
	(.146)		(.114)			(.091)		
Nonstockholders	**.057**		.017	.048		.049	.568	
	(.079)		(.070)			(.070)		
Bottom layer	.046		−.054	.570		.052	.828	
	(.186)		(.163)			(.158)		
Middle layer	.175		.350	.547		.173	.027	
	(.274)		(.207)			(.261)		
Top layer	.486		.417	.203		.292	.027	
	(.325)		(.235)			(.188)		
Nonstockholders vs. stockholders		3.255			4.340			3.296
		(.071)			(.037)			(.069)
Nonstockholders vs. top layer		2.146			2.378			2.941
		(.143)			(.123)			(.086)
	2. Single-Individual Households							
All	.202		−.008	.025		.261	.439	
	(.172)		(.106)			(.139)		
Stockholders	.698		.323	.005		.681	.969	
	(.496)		(.264)			(.326)		

TABLE 2
(*Continued*)

	Instrument Set 1		Instrument Set 2			Instrument Set 3		
	$\hat{\sigma}$	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ
Nonstockholders	.077		−.115	.071		.160	.250	
	(.143)		(.137)			(.123)		
Nonstockholders vs. stockholders		1.733			1.958			3.866
		(.188)			(.162)			(.049)
	B. Euler Equation for Treasury Bills 1. All Household Sizes							
All	.372		.362	.097		.264	.143	
	(.232)		(.225)			(.220)		
Bondholders	**.932**		.839	.215		.783	.147	
	(.368)		(.360)			(.353)		
Nonbondholders	**.105**		.087	.024		.005	.382	
	(.270)		(.257)			(.244)		
Bottom layer	.986		.798	.492		.726	.244	
	(.662)		(.612)			(.580)		
Middle layer	.287		.336	.476		.277	.627	
	(.550)		(.544)			(.523)		
Top layer	1.648		1.672	.356		1.530	.067	
	(.515)		(.505)			(.505)		
Nonbondholders vs. bondholders		4.029			3.192			4.127
		(.045)			(.074)			(.042)
Nonbondholders vs. top layer		8.055			8.905			7.871
		(.005)			(.003)			(.005)
	2. Single-Individual Households							
All	.681		.524	.032		.282	.053	
	(.456)		(.450)			(.411)		

Bondholders	2.624		2.759	.163		1.617	.007	
	(.885)		(.880)			(.824)		
Nonbondholders	.048		−.028	.039		−.105	.103	
	(.479)		(.475)			(.460)		
Nonbondholders vs. bondholders		8.371			7.527			4.722
		(.004)			(.006)			(.030)

NOTE.—Numbers in parentheses are standard errors for $\hat{\sigma}$ and p-values for the Wald test. For the overidentification test, the entries are p-values, and the test has two degrees of freedom for each of instrument sets 2 and 3. Twelve monthly dummies are included as explanatory variables and instruments. The estimations for all household sizes furthermore include $\Delta \ln$ (family size) as an explanatory variable and instrument. In addition the instrument sets include the following variables. Instrument set 1: log dividend-price ratio. Instrument set 2: log dividend-price ratio, lagged log real value-weighted NYSE return, and lagged log real Treasury bill return. Instrument set 3: log dividend-price ratio, default premium, and bond horizon premium.

not satisfy the Euler equation for the stock market index return to a reasonable approximation (e.g., because of transactions costs or poorly diversified equity portfolios). For nonstockholders, the estimate of the EIS is close to zero. A Wald test rejects that the EIS is the same for stockholders and nonstockholders at the 10 percent level. When adding more instruments, using instrument sets 2 and 3, one gets similar results, although the EIS estimates tend to decline a bit.

The estimation results for the Euler equation for the Treasury bill return are shown in panel B of table 2. They show the same patterns as the results based on the stock return, with a positive and significant estimate of the EIS around 0.8 for bondholders and a small, insignificant estimate for nonbondholders. Again, the EIS is estimated to be larger, around 1.6, for the top layer of bondholders.

The coefficients on $\Delta \ln$ (family size) (not shown) further support the hypothesis that the Euler equations for the stock return and the Treasury bill return hold for households that own the asset in question, but not for others. For those that hold the asset, the coefficient on $\Delta \ln$ (family size) is typically significantly different from zero with a point estimate around 0.6, whereas change in log family size has a smaller and insignificant effect in the Euler equations for households that do not hold the asset.

An important negative finding based on the instrumental variables estimations in table 2 is that there is no tendency for χ^2 tests of overidentifying restrictions to reject those restrictions for nonstockholders and nonbondholders but not for stockholders and bondholders, as would be expected. For instrument set 3 the overidentifying restrictions are rejected for stockholders and the overidentifying restrictions are rejected for nonstockholders and nonbondholders for instrument set 2. To determine the effects of restricting the EIS to be the same across the Euler equations for stocks and Treasury bills for a given group of households, table 3 shows the joint estimation of the two Euler equations for the stockholder/nonstockholder distinction. Under my classifications, all stockholders are also categorized as bondholders, and the tests of overidentification should reject the restrictions for nonstockholders but not for stockholders. As the table shows, the tests of overidentifying restrictions are now rejected for both stockholders and nonstockholders. The fact that the overidentifying restrictions are now consistently rejected even for stockholders is consistent with the previous results that EIS estimates for stockholders (around 0.3) are lower than EIS estimates for bondholders (around 0.8), despite the fact that most households in these two categories are the same according to my classifications. It may be the case that the much lower power of the instruments for predicting the real stock return leads to substantial small-sample bias of the instrumental variables estimator.

In addition to considering different instrument sets, I did several other robustness checks. First, the estimations were performed for households consisting of only one individual. The results, included in the tables, show even larger differences between stockholders and nonstockholders and between bondholders and nonbondholders. In my view, this most likely reflects the fact that these households face much simpler optimization problems, making it easier to detect the relation between consumption growth and asset returns. Second, I repeated the estimations assuming a representative agent within the set of stockholders and within the set of nonstockholders. The results (not shown) were similar. Third, I considered estimators that are invariant to whether consumption growth is regressed on asset returns or the other way around. For instrument set 1 this is not an issue since the model is just identified (when one Euler equation is estimated at a time); thus this leads to the same estimate of the EIS regardless of whether consumption growth or the asset return is used as a right-hand-side variable. For instrument sets 2 and 3, I reestimated the model using both the continuous-updating GMM estimator of Hansen, Heaton, and Yaron (1996) and a limited information maximum likelihood (LIML) estimator (for simplicity without corrections for heteroscedasticity and autocorrelation in the LIML estimation). For both of these estimators the EIS estimates were more robust across instrument sets than when the linear GMM estimator was used, with no tendency for the EIS estimates to decline when instruments are added. Finally, I considered an alternative to taking cross-sectional averages over stockholders and nonstockholders that should be more robust to measurement error in consumption in the case of a relatively small number of households in the cross section. If the Euler equation for a given asset return holds for each asset holder, one can use the stochastic discount factor based on the consumption of any of the asset holders to estimate it. As an alternative to the time series of cross-sectional averages of log consumption growth rates, I therefore constructed a times series in which the consumption growth rate for a given period for a given group of households was the median consumption growth rate for that period for that group of households. The corresponding time series of observations of $\Delta \ln$ (family size) for this particular set of households was constructed as well. Results corresponding to those in table 2 are shown in table 4. The results again confirm the importance of distinguishing between asset holders and non–asset holders. The estimates of the EIS for stockholders and for bondholders are larger than before, further supporting the limited participation hypothesis. The Euler equation for the stock return leads to EIS estimates around 0.4 for stockholders (larger for single stockholders), whereas the Euler equation for the Treasury bill return results in EIS estimates around one for bondholders.

TABLE 3
GMM ESTIMATION OF LOG-LINEARIZED EULER EQUATIONS: REAL TREASURY BILL RETURN AND REAL VALUE-WEIGHTED NYSE RETURN, JOINT ESTIMATIONS (CEX, 1982–96, Semiannual Data)

	INSTRUMENT SET 1			INSTRUMENT SET 2			INSTRUMENT SET 3		
	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ
	A. All Household Sizes								
All	.080	.025		.046	.036		.026	.016	
	(.077)			(.058)			(.045)		
Stockholders	**.442**	.022		.404	.071		.261	.011	
	(.144)			(.114)			(.079)		
Nonstockholders	**.013**	.036		.003	.041		.003	.028	
	(.057)			(.022)			(.033)		
Bottom layer	.006	.141		.003	.372		.007	.078	
	(.054)			(.021)			(.026)		
Middle layer	.068	.175		.242	.218		.048	.041	
	(.173)			(.203)			(.040)		
Top layer	.700	.020		.422	.190		.324	.004	
	(.354)			(.260)			(.179)		
Nonstockholders vs. stockholders			9.228			12.886			13.677
			(.002)			(.0003)			(.0002)
Nonstockholders vs. top layer of stockholders			3.890			2.127			4.211
			(.049)			(.145)			(.040)
	B. Single-Individual Households								
All	.187	.055		.024	.036		.141	.062	
	(.192)			(.041)			(.120)		
Stockholders	1.108	.028		.378	.023		.568	.043	
	(.527)			(.306)			(.278)		

Nonstockholders	.018	.114		−.028	.068		.050	.079	
	(.089)			(.050)			(.074)		
Nonstockholders vs. stockholders			4.161*			1.716*			3.250*
			(.041)			(.190)			(.071)

Note.—Numbers in parentheses are standard errors for $\hat{\sigma}$ and p-values for the Wald test. For the overidentification test, the entries are p-values, and the test has one degree of freedom for instrument set 1 and five degrees of freedom for instrument sets 2 and 3. Twelve monthly dummies are included as explanatory variables and instruments. The estimations for all household sizes furthermore include $\Delta \ln$ (family size) as an explanatory variable and instrument. In addition the instrument sets include the following variables. Instrument set 1: log dividend-price ratio. Instrument set 2: log dividend-price ratio, lagged log real value-weighted NYSE return, and lagged log real Treasury bill return. Instrument set 3: log dividend-price ratio, default premium, and bond horizon premium.

* These Wald tests do not allow for correlation of the error terms across the Euler equations for stocks and Treasury bills. If this were allowed, one of the matrices involved would not be positive definite.

TABLE 4
GMM Estimation of Log-Linearized Euler Equations: Real Treasury Bill Return and Real Value-Weighted NYSE Return, Separate Estimations (CEX, 1982–96, Semiannual Data, Median Household Approach)

	Instrument Set 1		Instrument Set 2			Instrument Set 3		
	$\hat{\sigma}$	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ
	A. Euler Equation for Stocks 1. All Household Sizes							
All	.125		.069	.016		.138	.691	
	(.081)		(.067)			(.064)		
Stockholders	**.437**		.451	.147		.287	.014	
	(.229)		(.175)			(.105)		
Nonstockholders	**.055**		−.014	.033		.113	.352	
	(.081)		(.073)			(.074)		
Bottom layer	.054		.082	.320		−.016	.408	
	(.231)		(.191)			(.201)		
Middle layer	−.082		.067	.097		.228	.006	
	(.340)		(.233)			(.278)		
Top layer	1.072		.933	.019		.706	.450	
	(.798)		(.538)			(.367)		
Nonstockholders vs. stockholders		2.833			4.900			4.203
		(.092)			(.027)			(.040)
Nonstockholders vs. top layer		1.687			2.331			4.116
		(.194)			(.127)			(.042)

	2. Single-Individual Households							
All	.348		.168	.028		.347	.877	
	(.257)		(.145)			(.182)		
Stockholders	1.152		1.030	.090		.874	.661	
	(.897)		(.622)			(.479)		
Nonstockholders	.238		.044	.058		.268	.624	
	(.206)		(.127)			(.153)		
Nonstockholders vs. stockholders		1.350			3.853			3.438
		(.245)			(.050)			(.064)
	B. Euler Equation for Treasury Bills 1. All Household Sizes							
All	.420		.413	.057		.231	.010	
	(.270)		(.262)			(.263)		
Bondholders	**1.380**		1.140	.189		.870	.047	
	(.507)		(.472)			(.425)		
Nonbondholders	**−.011**		−.101	.052		−.168	.153	
	(.290)		(.283)			(.268)		
Bottom layer	1.345		1.162	.060		1.214	.700	
	(.712)		(.677)			(.591)		
Middle layer	.267		.255	.295		.182	.075	
	(.691)		(.689)			(.660)		
Top layer	2.838		2.877	.346		2.390	.017	
	(.585)		(.565)			(.575)		
Nonbondholders vs. bondholders		7.476			8.733			7.077
		(.006)			(.003)			(.008)
Nonbondholders vs. top layer		22.514			24.930			20.094
		(.000)			(.000)			(.000)
	2. Single-Individual Households							
All	1.174		1.018	.155		.661	.073	
	(.504)		(.490)			(.433)		

TABLE 4
(*Continued*)

	Instrument Set 1		Instrument Set 2			Instrument Set 3		
	$\hat{\sigma}$	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ	$\hat{\sigma}$	Overidentification Test	Wald Test Equals σ
Bondholders	2.625		2.727	.253		2.256	.032	
	(.801)		(.794)			(.726)		
Nonbondholders	.740		.589	.085		.379	.047	
	(.617)		(.610)			(.564)		
Nonbondholders vs. bondholders		3.788			4.326			4.187
		(.052)			(.038)			(.041)

Note.—Numbers in parentheses are standard errors for $\hat{\sigma}$ and p-values for the Wald test. For the overidentification test, the entries are p-values, and the test has two degrees of freedom for each of instrument sets 2 and 3.. Twelve monthly dummies are included as explanatory variables and instruments. The estimations for all household sizes furthermore include $\Delta \ln$ (family size) as an explanatory variable and instrument. In addition the instrument sets include the following variables. Instrument set 1: log dividend-price ratio. Instrument set 2: log dividend-price ratio, lagged log real value-weighted NYSE return, and lagged log real Treasury bill return. Instrument set 3: log dividend-price ratio, default premium, and bond horizon premium.

Overall, the limited participation hypothesis is supported by the differences in coefficient estimates across stockholders and nonstockholders and across bondholders and nonbondholders, but not by the tests of overidentifying restrictions. The baseline results are robust to a variety of changes to the specification and estimation method. As discussed earlier, the more important specification choice is the use of semiannual data rather than higher-frequency consumption growth rates.

It should be pointed out that my estimations of the EIS are based on pretax returns since it is not possible to calculate accurate household-specific tax rates for capital income in the CEX.[7] It is, however, possible to get an idea of the bias taxes may introduce in EIS estimates. Suppose that the tax rate on the total net return (income return plus capital gains) to asset i is τ_i. Then the correct log return to holding asset i from t to $t+1$ is $\ln[1+(1-\tau_i)R_{it}]$, which fluctuates less than $\ln(1+R_{it})$. Thus ignoring taxes leads to a downward-biased estimate of the EIS since a smaller value is needed to "translate" the too large fluctuations in $\ln(1+R_{it})$ into consumption growth changes. With the approximation $\ln(1+x)\simeq x$, the right-hand-side variable should be $\ln[1+(1-\tau_i)R_{it}]\simeq(1-\tau_i)R_{it}$ but is $\ln(1+R_{it})\simeq R_{it}$. Thus I am approximately estimating $(1-\tau_i)\sigma$, not σ. This further strengthens the case that the EIS is not zero for stockholders and bondholders.

In comparison to the literature, it is relevant to discuss EIS estimates for households that hold savings accounts but not stocks, bonds, or mutual funds or U.S. savings bonds. The majority of households (around two-thirds in the CEX whether households are weighted using the CEX weights or not) have positive holdings in savings accounts. Since the return on such assets is fairly highly correlated with the Treasury bill return, one may have expected results based on aggregate consumption to lead to significantly positive estimates of the EIS. In the CEX, estimates of the EIS based on the Treasury bill return are still small, around 0.15, for households with savings account holdings that do not hold stocks, bonds, mutual funds, or U.S. savings bonds. They are around 0.4 but with a standard error around 0.7 using the real return on savings accounts, calculated on the basis of a series of passbook savings account returns obtained from the Federal Reserve Board. For both return series, the estimates for this group of households are somewhat sensitive to instrument choice. One possible explanation for mixed results for savings account holders is that most households with only savings (and typically also checking) account holdings have too little financial wealth

[7] As discussed earlier, it cannot be determined whether households reporting positive holdings of the category stock, bonds, or mutual funds hold both stocks and bonds or only one of these types of assets. When classifying households, I assumed that they hold both, but without specific information about whether this is correct, as well as dates of purchase, accurate household-level tax adjustments are infeasible.

for it to be optimal to do sophisticated intertemporal optimization. Another possibility is that households with small savings account holdings (and no stocks etc.) may in reality be net borrowers or be borrowing constrained, holding just a small amount of assets in a savings account for precautionary purposes. As mentioned in the Introduction, the findings of Gross and Souleles (2002) suggest that once a more appropriate interest rate is used for borrowers (in their case the credit card rate), borrowing and thus consumption growth are sensitive to the interest rate, suggesting a nonzero EIS.

III. Conclusion

The paper suggests that accounting for limited asset market participation is important for estimation of the elasticity of intertemporal substitution. Differences in the estimates of the EIS between asset holders and non–asset holders are large and statistically significant. This is the case whether one estimates the EIS on the basis of the Euler equation for the NYSE stock index return and distinguishes between stockholders and nonstockholders as best as possible or estimates the EIS on the basis of the Euler equation for Treasury bills and distinguishes between bondholders and nonbondholders. Estimates of the EIS are around 0.3–0.4 for stockholders and 0.8–1 for bondholders, and they are larger for households with larger asset holdings within these two groups.

I do not attempt here to explain why some households chose corner solutions for stocks or bonds. I have argued elsewhere (Vissing-Jørgensen 2002) that information costs work as entry costs and per period market participation costs, making it suboptimal for households with low or moderate financial wealth to enter these markets.

References

Attanasio, Orazio P.; Banks, James; and Tanner, Sarah." Asset Holding and Consumption Volatility." *J.P.E.* 110 (August 2002): 771–92.

Attanasio, Orazio P., and Browning, Martin. "Consumption over the Life Cycle and over the Business Cycle." *A.E.R.* 85 (December 1995): 1118–37.

Attanasio, Orazio P., and Weber, Guglielmo. "Intertemporal Substitution, Risk Aversion and the Euler Equation for Consumption." *Econ. J.* 99 (suppl., 1989): 59–73.

———. "Consumption Growth, the Interest Rate and Aggregation." *Rev. Econ. Studies* 60 (July 1993): 631–49.

———. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *J.P.E.* 103 (December 1995): 1121–57.

Brav, Alon; Constantinides, George M.; and Gezcy, Christopher C. "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence." *J.P.E.* 110 (August 2002): 793–824.

Brav, Alon, and Geczy, Christopher C. "An Empirical Resurrection of the Simple Consumption CAPM with Power Utility." Working paper. Chicago: Univ. Chicago, Grad. School Bus., 1996.

Daniel, Kent, and Marshall, David. "Equity-Premium and Risk-Free-Rate Puzzles at Long Horizons." *Macroeconomic Dynamics* 1, no. 2 (1997): 452–84.

Epstein, Larry G., and Zin, Stanley E. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57 (July 1989): 937–69.

Fama, Eugene F., and French, Kenneth R. "Business Conditions and Expected Returns on Stocks and Bonds." *J. Financial Econ.* 25 (November 1989): 23–49.

Gabaix, Xavier, and Laibson, David. "The 6D Bias and the Equity Premium Puzzle." In *NBER Macroeconomics Annual 2001,* edited by Ben S. Bernanke and Kenneth Rogoff. Cambridge, Mass.: MIT Press, 2002.

Gross, David B., and Souleles, Nicholas S. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Q.J.E.* 16 (February 2002): 149–85.

Hall, Robert E. "Intertemporal Substitution in Consumption." *J.P.E.* 96 (April 1988): 339–57.

Hansen, Lars Peter; Heaton, John; and Yaron, Amir. "Finite-Sample Properties of Some Alternative GMM Estimators." *J. Bus. and Econ. Statis.* 14 (July 1996): 262–80.

Hansen, Lars Peter, and Singleton, Kenneth J. "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *J.P.E.* 91 (April 1983): 249–65.

Ibbotson Associates. *Stocks, Bonds, Bills, and Inflation 1997 Yearbook.* Chicago: Ibbotson Assoc., 1997.

Jones, Larry E.; Manuelli, Rodolfo; and Siu, Henry. "Growth and Business Cycles." Working Paper no. 7633. Cambridge, Mass.: NBER, April 2000.

Lynch, Anthony W. "Decision Frequency and Synchronization across Agents: Implications for Aggregate Consumption and Equity Return." *J. Finance* 51 (September 1996): 1479–97.

Mankiw, N. Gregory, and Zeldes, Stephen P. "The Consumption of Stockholders and Nonstockholders." *J. Financial Econ.* 29 (March 1991): 97–112.

Parker, Jonathan A. "Spendthrift in America? On Two Decades of Decline in the U.S. Saving Rate." In *NBER Macroeconomics Annual 1999,* edited by Ben S. Bernanke and Julio J. Rotemberg. Cambridge, Mass.: MIT Press, 2000.

Vissing-Jørgensen, Annette. "Limited Stock Market Participation and the Equity Premium Puzzle." Working paper. Chicago: Univ. Chicago, Dept. Econ., 1999.

———. "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures." Working paper. Chicago: Univ. Chicago, Dept. Econ., 2002.

Zeldes, Stephen P. "Consumption and Liquidity Constraints: An Empirical Investigation." *J.P.E.* 97 (April 1989): 305–46.