

Asset Holding and Consumption Volatility

Orazio P. Attanasio

Institute for Fiscal Studies, University College London, and National Bureau of Economic Research

James Banks

Institute for Fiscal Studies and University College London

Sarah Tanner

Institute for Fiscal Studies

We investigate the possibility that limited participation in asset markets, and the stock market in particular, might explain the lack of correspondence between the sample moments of the intertemporal marginal rate of substitution and asset returns in U.K. data. We estimate ownership probabilities to separate “likely” shareholders from nonshareholders, enabling us to control for changing composition effects as well as selection into the group. We then construct estimates

This study has been financed under the Economic and Social Research Council Centre for the Microeconomic Analysis of Fiscal Policy at the Institute for Fiscal Studies. We are also grateful for funding from the European Union under the Training and Mobility of Researchers scheme, contract ERBFMRXCT960016. Material from the Family Expenditure Survey made available by the Office for National Statistics through the Economic and Social Research Council data archive has been used by permission of the controller of Her Majesty’s Stationery Office. Neither the ONS nor the ESRC data archive bears any responsibility for the analysis or interpretation of the data reported here. Thanks are due to Richard Blundell, Lucy Chennells, Lars Hansen, Costas Meghir, Annette Vissing-Jørgensen, Guglielmo Weber, and Frank Windmeijer for useful comments and discussions. We also received useful feedback from the participants at the February 1997 TMR meeting on savings and pensions, the summer 1997 NBER workshop on consumption, and the 1997 Econometric Society European meetings and at seminars at Bristol University, Oxford University, the Inter American Development Bank, the International Monetary Fund, Institut National de las Statistique et de Etudes Economiques, the International Institute at the University of Stockholm, and the University of Warwick. The usual disclaimer applies.

[*Journal of Political Economy*, 2002, vol. 110, no. 4]
© 2002 by The University of Chicago. All rights reserved. 0022-3808/2002/11004-0001\$10.00

of the IMRS for each of these different groups and consider their time-series properties. We find that the consumption growth of shareholders is more volatile than that of nonshareholders and more highly correlated with excess returns to shares. In particular, one cannot reject the predictions of the consumption capital asset pricing model for the group of households predicted to own both assets. This is in contrast to the failure of the model when estimated on data for all households.

I. Introduction

Asset pricing models based on the Euler equation for consumption have not performed well empirically. In these models, the pricing kernels are given by the intertemporal marginal rate of substitution (IMRS) of consumption. Aggregate consumption growth, which in standard representative agent models with standard preferences determines the marginal rate of substitution between present and future consumption, does not exhibit enough variability to be consistent with the observed time-series properties of asset prices and, in particular, with the mean and the variance of the excess return on shares over a relatively safe asset such as Treasury bills. But the equilibrium relationship between IMRS and asset returns holds only for individuals holding complete portfolios. As more detailed micro data on wealth and saving are made available, it is increasingly clear that the majority of individuals do not hold large stocks of financial wealth or fully diversified portfolios. This suggests that at least part of the equity premium puzzle discussed by Mehra and Prescott (1985) (see Kocherlakota [1996] for a survey) could be explained by the systematic difference of the time-series properties of consumption growth for asset market participants from those of aggregate consumption growth.

This was a point stressed by Mankiw and Zeldes (1991). They found that a distinction between shareholders and nonshareholders is important for resolving the empirical failure of consumption-based capital asset pricing models (CAPM). However, groups of shareholders and nonshareholders are defined on the basis of share ownership in the last period of the sample. Share ownership is neither a permanent nor an exogenous state of affairs. The time-series properties of the consumption growth of a group of individuals classified as share owners *at a single point in time* might not be indicative of the properties of the IMRS relevant for past asset prices. This is important in the United States and particularly in the United Kingdom, where levels of share ownership and the composition of the group of shareholders have changed dramatically in recent years. Second, Mankiw and Zeldes focus on the unconditional moments of consumption and asset returns. More efficiency

and power can be gained considering conditional moment restrictions. Third, the data they use contain information on food expenditure only. If food consumption is nonseparable from the other components of consumption, their estimates may be misleading.

This paper studies the time-series properties of shareholders' consumption and introduces a new way of controlling for the effects of compositional change. Panel data with a sufficiently long time-series dimension to allow us to identify groups of shareholders and nonshareholders over time, which also contain information on total consumption, do not exist. For this reason, we develop a grouping estimator for repeated cross-section data that conditions on *past* information to hold the composition of the group constant in looking at changes over time. This is an extension of synthetic panel estimation that, to our knowledge, has not been used before. We define groups of shareholders and nonshareholders in each time period on the basis of *predicted probabilities* of share ownership. Furthermore, we define these probabilities on the basis of variables that are perfectly predictable from one period to the next. In computing consumption growth, we compare the same group of households in adjacent periods; that is, we compute the IMRS between times t and $t + 1$ using the consumption of households predicted to be shareholders at time t . The lack of panel data forces us to use the technique we developed. However, we argue that our approach is better able to avoid the potential biases that might be induced by changes in the composition of stock ownership over time. Using the technique we develop, we estimate the Euler equations for share returns and bond returns imposing the cross-equation restriction on the parameter that represents the curvature of the utility function. We do not, however, model explicitly the conditional second moments of asset returns and consumption growth.

The data we use are drawn from the U.K. Family Expenditure Survey for the period 1978–95. This gives us a long and consistent time series of data on consumption, a crucial factor in estimating Euler equations and asset pricing models. Changing patterns of direct share ownership over the last 20 years also make the United Kingdom an interesting case for analysis. The Conservative government in the early 1980s introduced a number of measures designed to create a “share-owning democracy,” including the heavily advertised privatization of public utilities and tax breaks for employee share schemes. Largely as a result of these measures, there was a near trebling of the level of share ownership over a very concentrated period, 1985–88, which in all likelihood induced a change in the composition of shareholders. As we discuss below, changes in the composition of market participants can be quite problematic for the study of asset prices. The approach we propose is partially motivated by the need to deal with this problem, which we show to be important.

The paper is organized as follows. In Section II, we review briefly the consumption CAPM model and its testable implications for time-series properties of asset prices. Section III discusses in detail the econometric technique we develop to characterize the time-series properties of a variable for a group whose composition changes over time in a manner endogenous to the variable of interest. In Section IV, we present evidence on share ownership in the United Kingdom, and Section V uses the time-series properties of consumption growth of “likely” (defined by estimated probabilities of share ownership) and actual shareholders and nonshareholders to estimate the Euler equation for stocks and bonds and interprets the results. Section VI concludes the paper with some remarks about potential extensions.

II. The Consumption CAPM: Theory and Aggregate Evidence

A. Theory

Consider the standard intertemporal optimization problem facing a generic consumer with access to N different assets. Consumption and portfolio decisions are assumed to follow from the maximization of the expected lifetime value of utility from consumption (appropriately discounted) subject to an intertemporal budget constraint that reflects the intertemporal allocation possibilities available. Assuming that lifetime utility displays additive separability, we can write the maximization problem for household h as

$$\begin{aligned} & \max E_t \left[\sum_{s=t}^T U(c_{h,s}, v_{s,h}) \beta^{s-t} \right] \\ & \text{subject to} \\ & \sum_{k=0}^N A_{h,s}^k = \sum_{k=0}^N A_{h,s-1}^k (1 + r_s^k) + Y_{h,s-1} - C_{h,s-1}, \end{aligned} \quad (1)$$

where $C_{h,s}$ denotes (nondurable) consumption in period s ; $v_{h,s}$ denotes other factors, observable and unobservable, that might affect the (marginal) utility of nondurable consumption (such as demographic variables or taste shocks); β is the discount factor; $A_{h,s}^k$ is the amount of wealth held in asset k ; and r_s^k is the rate of return on that asset. If asset k is held in period t , a first-order condition for this problem is

$$\frac{\partial U(c_{h,t}, v_{h,t})}{\partial c_{h,t}} = E_t \left[\beta (1 + r_{t+1}^k) \frac{\partial U(c_{h,t+1}, v_{h,t+1})}{\partial c_{h,t+1}} \right], \quad (2)$$

and this holds for asset k independently of whether the consumer holds

other assets. Given the assumption of intertemporal separability, equation (2) can be rewritten as

$$E_i[m_{h,t+1}(1 + r_{t+1}^k)] = 1, \quad k = 1, \dots, n, \quad (3)$$

where $m_{h,t+1}$ is the IMRS between consumption in t and $t + 1$, and n are the assets held by the consumer in nonzero amounts, assumed to be, without loss of generality, the first n . For the remaining $N - n$ assets, equation (3) does not hold as an equality.

A key implication of consumption CAPM models, therefore, is that equilibrium returns are determined by a single factor: the IMRS. This fact gives rise to a number of orthogonality conditions, as is clear when one considers equation (3) for different assets. These orthogonality conditions can be used to estimate preference parameters in (1) and, provided that the model is overidentified, test the overidentifying restrictions. This was the approach followed, for instance, by Hansen and Singleton (1982, 1983), who estimated several versions of equation (3) using aggregate time-series data.

Equation (3) involves nonlinear relationships. As we are using a synthetic panel approach and we want to allow for measurement error, we prefer to deal with relationships that are linear in parameters. If the utility function is given by $U(c_{h,t}, v_{h,t}) = [c_{h,t}^{1-\gamma}/(1-\gamma)]e^{v_{h,t}}$ and we assume for notational simplicity that v is a scalar reflecting unobserved test shocks, we can log-linearize the Euler equation (3) for asset j to obtain

$$\ln(1 + r_{t+1}^j) = k_{h,t}^j + \gamma \ln \frac{c_{h,t+1}}{c_{h,t}} + \epsilon_{h,t+1}^j + v_{h,t+1} - v_{h,t} \quad j = 1, \dots, N, \quad (4)$$

where $\epsilon_{h,t+1}^j$ is an expectational error and, as mentioned above, v reflects unobserved heterogeneity; γ is the coefficient of relative risk aversion and, with expected utility and intertemporally separable preferences, the inverse of the elasticity of intertemporal substitution; and $k_{h,t}^j$ is a term including the log of the discount factor as well as conditional higher moments of the return on asset j and of consumption growth (such as variances and covariances). If one assumes lognormality of consumption growth and asset returns, $k_{h,t}^j$ is given by

$$k_{h,t}^j = -\ln \beta - \frac{\gamma}{2} \text{Var}_t(\Delta \ln c_{h,t+1}) - \text{Var}_t[\ln(1 + r_{t+1}^j)] \\ + \gamma \text{Cov}_t[\ln(1 + r_{t+1}^j), \Delta \ln(c_{h,t+1})], \quad (5)$$

where the subscripts t indicate that the variance and covariances are conditional on the information available at time t . If there are instruments that are uncorrelated with $\epsilon_{h,t}^j$, $\Delta v_{h,t+1}$ and with the innovations to $k_{h,t}^j$, the parameters of equation (4) can be estimated using generalized

method of moments (GMM) techniques.¹ Under the assumption of rational expectations and in the absence of measurement error, any variable dated $t - 1$ or earlier is a valid instrument.

In what follows we estimate a version of equation (4) for two different assets, imposing the cross-equation restriction that the coefficient on the rate of return is the same in the two equations. As the expression in equation (5) contains the variance of each rate of return and its covariance with consumption, we impose no restriction on the intercepts of the two equations. If one were able to model the conditional second moments, one could impose additional restrictions that could be informative about the ability of the model to explain observed excess returns.

B. *Asset Prices in the United Kingdom*

The main facts about the returns on shares and Treasury bills are not surprising. From the first quarter of 1978 to the last of 1995, as in the United States, the share returns are substantially higher and more volatile than the returns on Treasury bills. The average return on stocks is 0.0228 per quarter (standard deviation of 0.09) against a return of 0.0088 for Treasury bills (standard deviation of 0.012). These figures correspond to an annualized excess return of 0.0587 with a standard deviation of almost 0.3.

Estimating the Euler equations on these data and on aggregate consumption growth, one obtains the standard rejections of the overidentifying restrictions often reported in the literature. It has been shown (Attanasio and Weber 1993) that aggregation problems can introduce important biases in the estimation of Euler equations using aggregate data. In this context, the problems are exacerbated by the issue of nonparticipation. We therefore turn to the estimation of Euler equations on micro data that control for nonparticipation.

III. Methods

Our task is to estimate equation (4) for a group of stockholders. The first problem we have to face is the fact that our data are not a panel

¹ In a different approach, instead of considering a system of equations, Mankiw and Zeldes (1991) focused on excess returns and instead of conditional moments considered unconditional ones. They show that

$$\log \frac{E_t(1 + r_{t+1}^s)}{E_t(1 + r_{t+1}^b)} = \gamma \text{Corr}[\tilde{r}_{t+1}, \Delta \log(C_{t+1}) \sqrt{\text{Var}(\tilde{r}_{t+1}) \text{Var} \Delta \log(C_{t+1})}],$$

where $\tilde{r}_{t+1} = \log(1 + r_{t+1}^s) - \log(1 + r_{t+1}^b)$. As γ is the only parameter to be estimated, it can be identified by making the sample equivalents of the unconditional moments equivalent to the population moments.

but a (long) time series of cross sections. As different individuals are observed over time, we are forced to use a grouping estimator, of the kind proposed by Browning, Deaton, and Irish (1985) and used extensively in the literature. This is equivalent to aggregating equation (4) over the individuals belonging to a certain group *whose membership is fixed over time* and working with average (log) consumption for that group rather than individual consumption. Ignoring variations in participation, aggregating equation (4), we get

$$\ln(1 + r_{t+1}^j) = \bar{k}_t^j + \gamma(\overline{\ln c_{t+1}} - \overline{\ln c_t}) + \overline{\epsilon_{t+1}^j} + \overline{v_{t+1} - v_t}, \quad j = 1, \dots, N, \quad (6)$$

where the bar over a variable denotes the average over the group.

The second problem we have to tackle arises from the fact that the Euler equation for stocks between periods t and $t + 1$ holds only for households that hold stocks at t . Therefore, we want to look at the time-series properties of consumption growth for shareholders and non-shareholders (at time t) separately. In our time series of cross sections, we do not know whether a particular individual owning shares at a point in time owned shares in the previous quarter or will own shares in the following one. Therefore, we cannot use the standard synthetic panel techniques to follow the behavior of shareholders and nonshareholders. The basic problem is easily stated if we consider the quantity we would be using if we had a panel of stockholders. In this case, we could compute average consumption growth in equation (6) as

$$\begin{aligned} \Delta \ln(C_{t+1}^o) &\equiv [\overline{\ln(C_{t+1})} | d_t = 1] - [\overline{\ln(C_t)} | d_t = 1] \\ &= \frac{E_t \ln(C_{t+1} d_t)}{E_t[d_t]} - \frac{E_t \ln(C_t d_t)}{E_t[d_t]}, \end{aligned} \quad (7)$$

where d is a dummy indicating share ownership, and the operator E_t denotes the cross-sectional mean. With repeated cross sections we are unable to compute the first of the two terms on the right-hand side of equation (7).

One possibility would be to group individuals *at each point in time* on the basis of *current* share ownership. This procedure, however, is likely to give misleading results because of changes in the composition of the group of shareholders over time. These changes imply that the group of shareholders is constituted of different individuals at different points in time, so that changes in average consumption of a group so defined are not easily interpretable.

The third problem we have is that, even if we knew which individuals observed at t were holding stocks at t , if ownership changes over time, $\overline{v_{t+1} - v_t}$ would be computed over different groups of individuals over time. Exogenous changes in the incentives to participate in the stock market cause systematically different groups of consumers to enter the

market. Systematic differences in preferences among the stock owners at different points in time might introduce biases in the estimation of Euler equations since $\overline{v_{t+1}} - \overline{v_t}$ might be correlated with the instruments used in estimation.²

The approach we propose is to define groups in terms of their predicted ownership probabilities at a given point in time. We limit the variables that we use to predict ownership to those that do not vary over time or can be predicted perfectly, such as age (see Moffitt 1993). Because of this, given the estimated coefficients, we can, for households observed at time $t + 1$, compute the probability of ownership at time t . For each pair of adjacent time periods, we define groups of likely shareholders and nonshareholders according to their ownership probability in the first of the two periods and compute the consumption growth for these groups.

To be more precise, we define the consumption growth of what we call the group of “likely shareholders” as

$$\begin{aligned} \Delta \ln(C_{t+1}^s) \equiv & [\overline{\ln(C_{t+1})} | p(\text{shareholder})_t > p_t] \\ & - [\overline{\ln(C_t)} | p(\text{shareholder})_t > p_t], \end{aligned} \quad (8)$$

where the notation $[\overline{x_t} | A]$ denotes the sample mean of the variable x for the individuals observed at time t for whom condition A is satisfied. The term $p(\text{shareholder})_t$ is the predicted probability of owning shares, and p_t is a cutoff point. Notice that the condition in the two terms on the right-hand side of (8) is the same. We compute a similar expression for the “unlikely shareholders.”

In our empirical analysis, we experiment with two alternatives for the cutoff point: the actual proportion of shareholders in each period t of our sample and an arbitrary cutoff point that is constant over time. We discuss the motivation for these choices below. What is important, however, is that the groups defined in each pair of subsequent time periods t and $t + 1$ are formed on the basis of the same criterion: the probability of ownership at time t . For this reason it is crucial that the variables we use to predict ownership are fully predictable or constant over time. To compute the time-series properties (variability, correlation with expected risk premium, and so on) of the consumption growth of likely shareholders, we compute the averages in equation (8) for all pairs of adjacent time periods in our sample.

² Share ownership is likely to be endogenous with respect to consumption. When hit by a shock, a household might decide to invest in stocks when it did not before (if the shock is positive) or might decide to liquidate its holding of stocks (if the shock is negative). However, if the differences among households in and out of the market are not systematic and if these shocks are unanticipated, this would not bias the results of a properly aggregated Euler equation.

Given our discussion above, it should be clear that our procedure might yield inconsistent estimates of the parameters of interest for two reasons, arising from the lack of a longitudinal dimension in our data and from the presence of unobserved heterogeneity, respectively. We discuss them in turn. Note that the unobserved heterogeneity might induce biases even if we had panel data on share ownership and consumption. We claim below that our approach is better suited to deal with this problem.

First, the expression for the consumption growth of likely shareholders on the right-hand side of equation (7) differs from the expression on the right-hand side of equation (8) by four terms reflecting prediction errors from the probit. Two possible misclassifications can occur. Individuals observed either at t or at $t + 1$ who hold shares can have $p(\cdot) < p_b$ and hence their consumption is not counted in our measure of the IMRS. Similarly, individuals with $p(\cdot) < p_t$ who do not hold shares will have their consumption falsely included.

If we could predict share ownership perfectly, then $\Delta \ln C_{t+1}^s = \Delta \ln C_{t+1}^o$, and we would be measuring the consumption growth of actual shareholders at time t . When this is not the case, the issue is whether the classification errors induced by the estimation introduce a bias in the estimates of the structural parameters and in the tests of overidentifying restrictions. This in turn depends on the covariance between these (unobserved) terms and the instruments used in our GMM procedure. If the share ownership equation includes time dummies (or is a very flexible function of time trends), the classification errors are likely to be uncorrelated with the aggregate instruments we use (such as lagged interest rates and so on). Moreover, as we use instruments that are lagged two periods, our procedure should still yield consistent estimates unless there are reasons to believe that the classification errors exhibit serial correlation.³

Second, while our approach defines groups consistently for any two subsequent periods, so that the definition of “consumption growth” makes sense, group membership changes when one considers different observations (over time) for the rate of consumption growth. The composition of the group of likely (and actual) shareholders changes over time both because the probability of ownership may change over time and because the cutoff point might change. If utility depended only on consumption and not on unobserved taste shocks or if these taste shocks were constant over time, this would not constitute a problem. However, as discussed above, systematic changes in the composition of the group

³ A reductive interpretation of our procedure would see it as a test of the null hypothesis that limited stock ownership is not the explanation of the empirical failure of the consumption CAPM. That is, we check whether the consumption behavior of what we define as the likely shareholders is systematically different from that of the rest of the population.

might induce changes over time in $\overline{v_{t+1} - v_t}$ that might be correlated with the instruments used in estimation and therefore imply a bias. Notice that this problem is not related to the lack of a longitudinal dimension of our data. It would be a problem even if we were to know who among the individuals interviewed at $t + 1$ held stocks at t .

Our different assumptions on the cutoff point are meant to deal with this problem. By choosing a fixed cutoff, we are likely to minimize the biases induced by changes in the composition of the group. By choosing a time-varying cutoff point, we are likely to minimize the biases induced by misclassification.

IV. Patterns of Share Ownership in the United Kingdom

The overwhelming majority of U.K. households—more than 75 percent—do not own shares directly. Furthermore, levels of share ownership greater than 20 percent are a recent phenomenon. The level of direct share ownership changed dramatically in the United Kingdom in the 1980s, more than trebling over the period 1985–88.⁴ This is shown in figure 1 with data from the Family Expenditure Survey (FES). Although the FES contains little information on the amounts of assets held by households, data on dividend and interest income can be used to infer ownership of different assets, including stocks and shares. This procedure for identifying share owners yields levels of ownership that are very similar to those found in other data sources for individual years.

The rapid increase in share ownership coincided with a number of measures designed to promote a “share-owning democracy.” The first was a program of privatization and the heavily advertised flotation of a number of public utilities including British Telecom (1984) and British Gas (1986). Privatization accounts for a large part of the increase in the number of share owners. More detailed information on share ownership contained in the 1987 and 1988 General Household Surveys, for example, shows that more than half of all shareholders owned shares in privatized companies and that a large proportion of them owned shares in only a privatized company. However, the evidence also suggests that there was a more general increase in ownership of shares. The information in the General Household Surveys shows that there was an increase in the proportion of people owning “other” (i.e., nonprivatized)

⁴ At the same time as the proportion of the population owning shares directly has increased, the total proportion of shares owned directly by the personal sector has fallen. In 1957, nearly two-thirds of all shares were owned directly by individuals. By 1975 the figure was 37.5 percent, and by 1994 it had fallen to 20.3 percent. This is largely the result of a rapid growth in institutional ownership by pension funds and insurance companies. In this paper the group of shareholders is defined only as those who hold shares directly since the equilibrium relationships described in the previous section should hold for this subset of individuals.

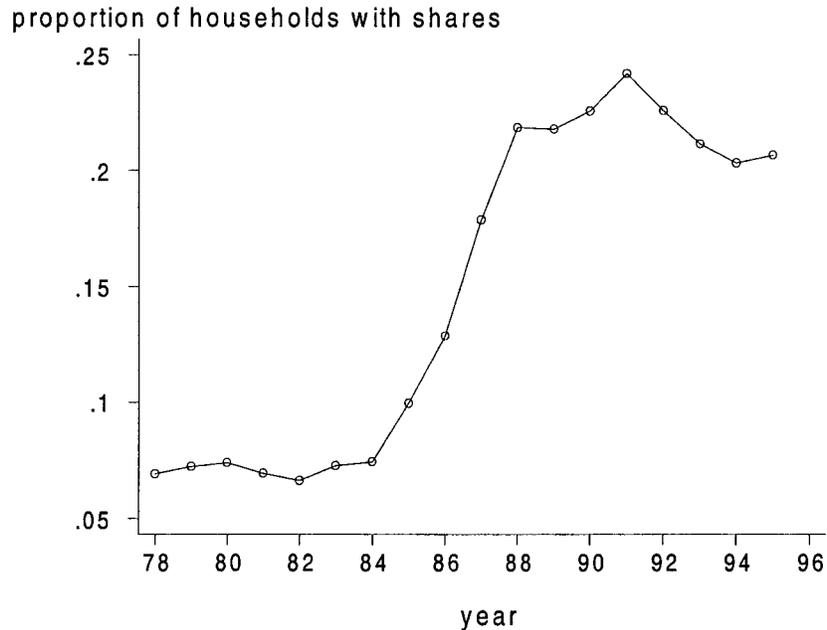


FIG. 1.—Share ownership, 1978–95

shares from 10 percent in 1987 to 13 percent in 1988. This may have been partially brought about indirectly through privatization and the knock-on effects of increased awareness of and knowledge about share ownership.⁵ A second measure introduced by the government was tax-favored employee share schemes, three of which were brought in between 1979 and 1984. By 1988 the total number of employees participating in such schemes was over 1.1 million. More recently, in 1988 the government introduced tax-free personal equity plans for limited direct and indirect equity holdings, and by 1999 these plans were held by 10 percent of the population. Toward the end of the 1990s the conversion of mutual building societies into banks also resulted in windfall allocations of shares to millions of building society account holders and mortgagors.

As the level of share ownership increased during the 1990s, the com-

⁵ Haliassos and Bertaut (1995) explain the low levels of direct share ownership by appealing to inertia and lack of information. Their conclusion is that an increase in share ownership may be brought about by extensive initial advertising plus a continuous flow of information, but this may not be effective in drawing stockholders from lower-income groups.

TABLE 1
PROBIT ESTIMATION FOR SHARE OWNERSHIP

	Parameter	Standard Error	Marginal Effect
Age of head	.281	.022	.055
(Age of head) ²	-.021	.006	-.004
Head has A levels	.697	.031	.160
College education	1.198	.039	.349
Age × A levels	.082	.013	.016
Age × college	.117	.016	.023
Trend	3.522	.875	.687
Trend ²	-31.553	5.060	-6.158
Trend ³	104.932	12.181	20.479
Trend ⁴	-127.493	12.888	-24.882
Trend ⁵	51.629	4.960	10.076
Age × trend	-.212	.083	-.042
Age × trend ²	-.365	.047	-.071
A levels × trend	-.634	.058	-.124
College × trend	.122	.076	.024
Constant	-2.022	.051	...
Observations	83,736		
Pseudo R ²	.1203		

position of the group of shareholders also changed, which suggests that controlling for the changing composition of shareholders over this period may be an important issue. In general, shareholders tend to be older and better educated than the rest of the population, but these differences have been getting smaller over time. Descriptive evidence from the data shows that the average age of those owning stocks and shares has fallen from 56.5 in 1978 to 51.7 in 1995. The differential associated with higher levels of education has also fallen over time. In 1978, 63.7 percent of households with shares had a head with postcompulsory education, compared with 33.5 percent of all households. By 1988, the proportion of share-owning households with heads with postcompulsory education had fallen to 61.7 percent, whereas the proportion of all household heads with postcompulsory education had actually increased to 41.3 percent.

These effects can be seen clearly in the results of the probit regression we use to predict share ownership, which is presented in table 1. The probit is estimated on pooled data from the FES for 1978–95—more than 80,000 households in total. For the reasons outlined in the previous section, the right-hand-side variables are constrained to be those that are constant over time or those that evolve in a predictable way. From the set of potential demographic and economic variables in the FES, we select age, education, and time. However, we adopt a very flexible specification including higher-order terms in age and time. We also interact the time trend with the other explanatory variables to allow for the fact that the effects of age and education appear to change over

time. We obtain very similar results by estimating a different probit for each year in the sample. The results show that the probability of share ownership increases with age, time, and higher levels of education, although the positive effects of college education and A levels on share ownership diminish over time.

V. Results

We start our exercise by simply comparing the time-series properties of consumption growth for the two groups of predicted (or likely) shareholders and nonshareholders. On average, the predicted shareholders have higher consumption growth over the period than nonshareholders: 0.29 percent per quarter for shareholders compared to 0.06 percent for nonshareholders when we use the variable cutoff point and 0.12 and 0.08 when we use the fixed one. The standard deviation of the time series of average consumption growth (adjusted to take account of the variance induced by sampling error and differences in cell size)⁶ is almost twice as high: 3.7 for shareholders compared to 2.5 for nonshareholders when we use a variable cutoff point and 3.3 compared to 2.6 when we use a fixed cutoff. If we compute the rate of growth for actual share owners and non-share owners, the difference is even more dramatic. The standard deviation of the time series of average consumption growth (again adjusted for sampling error) is almost twice as high: 6.52 for shareholders compared to 3.37 for nonshareholders.

These numbers compare to a standard deviation of 0.88 percent for per capita quarterly consumption growth in the National Income and Product Account series over the same time period. Some of this dramatic difference may arise from inconsistencies in the two series definitions. In particular, there are differences in the included consumption items, in the sector definitions (namely, the presence of nonhouseholds in the national accounts series), and in the per capita versus per adult equivalent conversion of total expenditure. However, the most important differences are likely to be due to aggregation issues, consistent with the results for the United States of Vissing-Jørgensen (1999), who points out that aggregation bias can be substantial in computing the mean and standard deviation of the IMRS.

This evidence, consistent with that presented by Mankiw and Zeldes (1991), suggests that differences between share owners and non-share owners are likely to be important in estimating consumption growth

⁶ This adjustment is important: the fact that the sample of stock owners is smaller than the sample of non-stock owners implies that sampling variability accounts for part of the difference in the time-series variance of consumption growth of the two groups. Fortunately, we have information on sample sizes and within-cell variability that allows us to correct for this effect.

and intertemporal substitute parameters. But a greater variance is a necessary but not a sufficient condition for resolving the puzzle. One would observe higher variances if the sample of stockholders were more affected by measurement error. In conjunction with a higher variance, we need a higher (conditional) correlation of the IMRS of shareholders with the excess returns on shares. It is hard to imagine how measurement error would induce such a correlation with intertemporal prices. For this reason we turn to the estimation of Euler equations. In particular, we estimate several versions of the Euler equation given by expression (4) that differ for the sample used to compute average consumption growth: the entire sample, actual shareholders, and likely and unlikely shareholders. More precisely, we consider the return on stocks (r_t^s) and Treasury bills (r_t^b) and estimate, for each of the samples of likely and unlikely share owners, the two Euler equations:

$$r_t^s = \gamma_1 \Delta \ln c_t + \beta_1 + u_t^s, \quad (9a)$$

$$r_t^b = \gamma_2 \Delta \ln c_t + \beta_2 + u_t^b. \quad (9b)$$

To estimate the parameters of these equations, we use the normalization-free GMM procedure proposed by Hansen and Singleton (1996). However, given the nature of our data and problem, we have to modify it slightly. First, unlike Hansen and Singleton, we do not use only lagged interest rates and consumption growth as instruments. Second, since the residuals of our expression are characterized by an MA(1) process, both because of the limited sample size at the micro level and because of potential time aggregation effects, we use instruments dated $t - 2$ or earlier. Finally, to improve efficiency, we take into account explicitly the nature of the residuals to construct a feasible generalized least squares type of estimator and (forward) filter the consumption growth and interest rate data after obtaining an estimate of the MA(1) parameter in a first step. Details of the estimation procedure are given in Appendix B. Consumption growth is measured as the change in the log of deseasonalized consumption per adult equivalent in the various groups considered. The instruments, which include the second lag of consumption growth and several financial variables, are listed in the notes to tables 2 and 3. In each table, we also report the R^2 for the reduced-form regressions for $\Delta \ln c_t$, r_t^b , and r_t^s . In columns 1, 2, 4, and 5, we report the point estimates of the intercepts and slopes of equations (9a) and (9b) (β 's and γ 's). For each sample we consider, we estimate the system both with and without the cross-equation restriction that $\gamma_1 = \gamma_2$. Columns 3 and 6 contain a test of the overidentifying restrictions, along with its p -value in parentheses.

First, we present the results obtained in our sample if one uses actual

TABLE 2
GMM ESTIMATION OF EULER EQUATIONS FOR RETURNS ON SHARES AND TREASURY BILLS

	SHARE EQUATION			TREASURY BILL EQUATION		
	β_1 (1)	γ_1 (2)	J (3)	β_2 (4)	γ_2 (5)	J (6)
A. Whole Sample						
Unrestricted	.022 (.030)	6.853 (2.779)	5.85 (.441)	.009 (.006)	1.448 (.513)	7.362 (.289)
Restricted	.026 (.021)	5.226 (1.492)005 (.020)	5.226 (1.492)	14.43 (.344)
First-stage R^2 : $\Delta \ln c_t$: .045; Δr_t^b : .484; Δr_t^s : .171						
B. Shareholders						
Unrestricted	.029 (.030)	3.527 (.916)	10.165 (.118)	.009 (.010)	.548 (.149)	9.34 (.155)
Restricted	.034 (.013)	1.439 (.337)010 (.010)	1.439 (.337)	18.244 (.148)
First-stage R^2 : $\Delta \ln c_t$: .082; Δr_t^b : .469; Δr_t^s : .10						
C. Nonshareholders						
Unrestricted	.025 (.064)	16.442 (7.782)	4.055 (.669)	.009 (.033)	8.564 (4.165)	5.667 (.462)
Restricted	.036 (.067)	-18.107 (7.181)013 (.066)	-18.107 (7.181)	13.468 (.412)
First-stage R^2 : $\Delta \ln c_t$: .020; Δr_t^b : .495; Δr_t^s : .167						

NOTE.—Standard errors are in parentheses. Instruments include a constant, a dummy for financial liberalization, and the second lag of the rate of growth in consumption, share returns, Treasury bill returns, the spread between three-month Treasury bills and 20-year bonds, the spread between 20-year and five-year bonds, and inflation. The three R^2 values refer to those of the first step for consumption growth, return on Treasury bills, and return on stocks, respectively, in each of the samples. The sample period is 1978:4–1995:4. In cols. 3 and 6, J is a test of the overidentifying restrictions.

share ownership rates to define the two sample groups as opposed to predicted or likely ownership. In the three panels of table 2, we report estimates for the whole sample and for the subsample of actual shareholders and nonshareholders. As mentioned above, the owners at times t and $t + 1$ are not necessarily the same, which constitutes a problem for the interpretation of consumption changes from one period to the next.

In panel A, the estimates for the total sample indicate a very high 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 Treasury bills jointly. On the other hand, we obtain a much smaller value from the equation on Treasury bills only. In neither case, however, does the test of overidentifying restrictions reject the null. Panels B and C show that although the estimates are somewhat better for shareholders, the same is not true for nonshareholders. In particular, the restricted estimate of the coefficient of relative risk aversion is 1.439 in the sample of shareholders and -18.1 in the sample of nonshareholders. Notice that for the group of shareholders, however, the test of overidentifying

TABLE 3
GMM ESTIMATION OF EULER EQUATIONS: PREDICTED SHAREHOLDERS

	SHARE EQUATION			TREASURY BILL EQUATION		
	β_1 (1)	γ_1 (2)	J (3)	β_2 (4)	γ_2 (5)	J (6)
A. Predicted Shareholders (Variable Cutoff)						
Unrestricted	.0170 (.015)	1.498 (.529)	1.550 (.956)	.006 (.008)	1.027 (.365)	8.231 (.222)
Restricted	.030 (.010)	.647 (.180)007 (.005)	.647 (.180)	13.47 (.415)
First-stage R^2 : $\Delta \ln c_t$: .118; Δr_t^b : .473; Δr_t^s : .178						
B. Predicted Nonshareholders (Variable Cutoff)						
Unrestricted	.0207 (1.592)	-4.861 (2.093)	5.976 (.426)	.009 (.027)	7.179 (2.684)	8.724 (.190)
Restricted	.0492 (.097)	-27.058 (8.440)0288 (.096)	-27.058 (8.440)	14.58 (.334)
First-stage R^2 : $\Delta \ln c_t$: .038; Δr_t^b : .504; Δr_t^s : .164						
C. Predicted Shareholders (Fixed Cutoff)						
Unrestricted	.0180 (.0233)	3.175 (1.014)	5.22 (.516)	.007 (.010)	1.366 (.467)	9.409 (.152)
Restricted	.032 (.011)	1.065 (.325)007 (.007)	1.065 (.325)	13.22 (.431)
First-stage R^2 : $\Delta \ln c_t$: .172; Δr_t^b : .471; Δr_t^s : .171						
D. Predicted Nonshareholders (Fixed Cutoff)						
Unrestricted	.0216 (.051)	12.868 (4.824)	6.279 (.393)	.009 (.008)	2.121 (.799)	7.826 (.251)
Restricted	.015 (.075)	20.524 (6.503)	...	-.005 (.075)	20.524 (6.503)	14.30 (.353)
First-stage R^2 : $\Delta \ln c_t$: .031; Δr_t^b : .506; Δr_t^s : .161						

NOTE.—Standard errors are in parentheses. Instruments include a constant, a dummy for financial liberalization, and the second lag of the rate of growth in consumption, share returns, Treasury bill returns, the spread between three-month Treasury bills and 20-year bonds, the spread between 20-year and five-year bonds, and inflation. The three R^2 values refer to those of the first step for consumption growth, return on Treasury bills, and return on stocks, respectively, in each of the samples. The sample period is 1978:4–1995:4. In cols. 3 and 6, J is a test of the overidentifying restrictions.

restrictions is close to rejecting the null (p -value of .1), probably as a consequence of the large differences between the unrestricted slope estimates in the two equations. Moreover, the estimates of γ obtained from the Euler equation for Treasury bills vary considerably across samples.

As mentioned above, changes in ownership patterns would cause these results to be biased. For this reason we present in table 3 the results we obtain using our approach, which uses predicted share ownership status to define groups that are constant over adjacent periods. As we discussed in Section III, there are several ways in which we can define the group of the likely shareholders, depending on the definition of the cutoff point p_t in equation (8). In estimating Euler equations such as (4), we

experiment with two alternatives. First, we choose the point p_i to coincide with the average proportion of shareholders in each period of our sample. Second, we keep p_i fixed at 20 percent in all sample periods. As discussed above, the first definition is likely to minimize the “misclassification” of shareholders, whereas the second is aimed at minimizing the problems induced by changes in the composition of the group of shareholders.

Panels A and C of table 3 give the results for the predicted shareholders (defined according to the variable and fixed cutoffs, respectively) and show that the model performs well for this group. More specifically, in both cases the estimate of the risk aversion coefficient is plausible and is estimated with reasonable precision. Moreover, the overidentifying restrictions are not rejected. The point estimate of γ is very close to unity in the sample of likely shareholders defined with the fixed cutoff and a bit smaller in the sample with the variable cutoff. In panels B and D, we report the results for the likely nonshareholders. These estimates indicate very implausible values for the coefficient of risk aversion when this is estimated either with the share returns equation or with the two equations jointly, as would be expected, since the Euler equation for shares should not hold for this group. In the variable cutoff sample, the estimates are negative; in the other sample, they are positive but implausibly large. Although the overidentification restrictions are not rejected, the low R^2 for consumption growth suggests that this may be due to weak instruments.⁷

The estimate of γ obtained when one uses the equation for Treasury bills only in the fixed cutoff point is not significantly different across groups. In the sample with the variable cutoff point, however, the point estimate from the Treasury bill equation is considerably larger (even though the parameter is estimated with low precision) in the group of unlikely shareholders. If we think that most people in the sample hold some kind of safe asset, this evidence seems to indicate that the biases induced by the changes in composition of shareholders (more relevant for the variable cutoff point sample) are more important than those caused by the misclassification (which are more likely to be relevant in the fixed cutoff point sample). For this reason, our preferred set of estimates are those in panel A. Values of γ just above unity are not inconsistent with the estimates of the elasticity of intertemporal substitution obtained from the FES data on different sample periods by At-

⁷ In a previous version of the paper we showed that a Hayashi-Sims type GMM estimator that was not normalization-free gives estimates of the structural parameters similar to those reported, but also that the test of overidentifying restrictions rejects the null for the sample of the likely nonshareholders and for the whole sample and does not reject it for the sample of likely shareholders.

tanasio and Weber (1993) and Blundell, Browning, and Meghir (1994), among others.

Notice that the point estimates of γ obtained from the restricted system in the sample of likely shareholders are lower but not too different from those obtained with the sample of actual shareholders, especially with the fixed cutoff. However, when actual shareholders are used, unlike the results for the fixed cutoff point in table 3, the point estimates of γ obtained from only the Treasury bill equation change considerably across groups, once again pointing to the importance of the biases induced by changes in composition.

The interpretation of the constants of equations (9a) and (9b) is not completely straightforward since these constants include the discount factors and the unconditional second moments of asset returns and consumption growth, as shown in equation (5). However, the difference between the two constants should be equal to the difference between the variances of the two asset returns and the difference between the covariances of each of them with consumption. If we take the unconditional variance of the two assets to be as in the sample period (0.09^2 and 0.012^2) and $\gamma = 1.065$ (as estimated in panel C), the implied difference between the covariance of consumption with share returns and that with Treasury bill returns would be about 0.026 (if one uses the point estimates for the β 's from the restricted system) or 0.019 (if one uses those from the unrestricted one). This is too high, in absolute value, to be consistent with the observed variances of consumption growth and asset returns, but given the lack of precision with which the constant terms are estimated, such an inference from these point estimates is not particularly appropriate.

VI. Conclusions

This paper has looked at the empirical failure of the consumption asset pricing model in the context of recent secular changes in the number and type of shareholders in the United Kingdom. Since the first-order conditions for the model hold as an equality only for individuals that are currently participating in asset markets, it is natural to look at the consumption behavior of these individuals rather than the aggregate population. Pursuing this empirical strategy poses a number of problems. Not only do we need household-level information on consumption and on asset ownership, but we also have to deal with the fact that asset ownership is neither a permanent nor an exogenous status for the households in the survey. In addition, the available data for the United Kingdom, while providing excellent information on consumption and share ownership, are not a panel. This is problematic since the Euler equation holds only for those owning shares in adjacent periods. To deal with

this, we develop an extension of the synthetic cohort technique that defines groups of individuals with constant membership at adjacent dates on the basis of the estimated probabilities of owning stocks. In the empirical application, therefore, we define a group of “likely shareholders.”

We obtain strong results. When we estimate Euler equations for the group of likely shareholders, using both a single asset and two assets simultaneously, we obtain sensible values for the parameter of interest (the elasticity of intertemporal substitution) and we fail to reject the overidentifying restrictions implied by the model. For the other groups (the total sample and the unlikely shareholders), we obtain unappealing estimates of the structural parameters. This last result is important in showing that there is some empirical power in our approach. We also present some evidence that indicates that the bias induced by changes in the composition of share owners is important: estimates of the elasticity of intertemporal substitution vary widely across groups when we divide the sample on the basis of actual ownership and are much more stable when we use the method we devise to minimize this type of bias.

In summary, we have shown that the time-series properties of shareholders’ consumption growth are very different from those of non-shareholders and aggregate consumption growth. And they are different in a way that is consistent with some of the implications of the consumption CAPM. In particular, estimates of the elasticity of intertemporal substitution obtained from two Euler equations for shareholders (one for bond returns and one for share returns) are consistent with each other and theoretically plausible. We have not tried to model the conditional second moments, whose means enter the intercept of the Euler equation we have estimated. These are important restrictions that should be checked to claim whether the so-called equity premium puzzle can be explained by limited participation and differences in the time-series properties of the consumption of stock owners and non-stock owners.

Our results suggest a number of extensions. Given the limitations in our sample period and the reliance on time-series variation to identify the parameters of interest, we have worked with very simple preference specifications. It would be interesting to work with preferences that are more general in the way demographic and labor supply factors are allowed to affect utilities. More generally, it would also be interesting to consider more flexible forms of preferences, including the non-expected utility preferences of the kind studied by Epstein and Zin (1989) and models with habit formation. In particular, Attanasio and Weber (1989) showed that with Epstein-Zin preferences and an unobservable market portfolio, in a system like (9), γ represents the inverse of the elasticity of intertemporal substitution. The coefficient of relative

risk aversion is buried in the intercepts of the system and is not easily identifiable. These preferences might be appealed to in order to justify the fact that, as we discussed at the end of Section V, the point estimates of our intercepts seem inconsistent with the relative magnitude of the unconditional variances of returns.

The most important challenge, however, and the puzzle that our study leaves unresolved, is to explain the limited ownership of shares more structurally, particularly given the size of average excess returns. The descriptive evidence we present is suggestive and shows that the increase in ownership was quite widespread in the population. While it was probably triggered by the privatizations of the mid 1980s and by the associated publicity, the trends cannot be explained only by that episode or by the ownership of shares in privatized firms. Understanding the factors, such as fixed costs, that still prevent ownership for large sectors of the population remains an important topic for future research.

Appendix A

The Data

To estimate our model we make use of 18 consecutive years of the Family Expenditure Survey running from 1978 to 1995. The FES is collected by the Office for National Statistics, primarily with the purpose of computing weights for the retail price index. However, the survey contains detailed information on the income and demographic characteristics of British households as well as their expenditure patterns. For our analysis we select only households in which the head of the household is between ages 20 and 59. This leaves a total of 83,736 households, an average of 1,163 per quarter (with a minimum of 1,008 and a maximum of 1,269). As discussed in the text, share ownership is imputed from responses to the income part of the questionnaire, which collects information on detailed components of income. In the case of asset-related components, the questions refer to the previous 12 months. The proportion of imputed shareholders observed in each quarter has a minimum of 0.0627 and a maximum of 0.2621 over our sample period. The time-series pattern of this proportion corresponds closely to other sources of data on direct ownership of shares in the periods in which they are available.

For the purposes of predicting likely share ownership, we use a specification interacting age, education, and time effects. Age is defined as the current age of the head of the household only. A similar criterion applies to education, where we construct three measures of household education according to whether the head of the household left full-time education at or before the compulsory school-leaving age, between this age and 18, and above age 18. We refer to the penultimate group as having A levels and the final group as having college education.

Information on expenditure is collected by means of a two-week diary covering all expenditures by all household members. We construct a definition of total nondurable expenditures that excludes expenditures on housing, cars, electrical equipment, furniture, and furnishings. Nondurable expenditures are averaged over the two-week period, aggregated to the household level, and then divided

by an equivalence scale (to take account of the composition of the household) and deseasonalized for use in our analysis. A detailed analysis of FES income and spending data over time shows that, for nondurable items in particular, the FES captures a large and constant proportion of household expenditures as measured from other sources of aggregate information (see Banks and Johnson [1998] for further details of these results as well as a detailed description of the FES sample).

Appendix B

Estimation Methods

In this Appendix we describe in detail the GMM procedure we used to estimate the Euler equations in Section V. As mentioned, we modify the normalization-free GMM procedure proposed by Hansen and Singleton (1996). The essence of the method, which makes it normalization-free, consists in noticing that given a value of γ , the reduced-form parameters can be written in closed form in a very simple way. In particular, given a $T \times k$ matrix of instruments \mathbf{z} (which does *not* include a constant term), the three endogenous variables $\Delta \ln c$, r^s , and r^b , and a value of γ , one can write a reduced form consistent with the structural system in (9) as

$$\begin{pmatrix} \Delta \ln c \\ r^b \\ r^s \end{pmatrix} = \begin{pmatrix} \mathbf{i} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{i} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{i} \end{pmatrix} \begin{pmatrix} \pi_c \\ \pi_b \\ \pi_s \end{pmatrix} + \begin{pmatrix} \mathbf{z} \\ \gamma \mathbf{z} \\ \gamma \mathbf{z}' \end{pmatrix} \boldsymbol{\Pi} + \begin{pmatrix} \mathbf{u}_c \\ \mathbf{u}_b \\ \mathbf{u}_s \end{pmatrix}, \tag{A1}$$

where the \mathbf{i} 's, the $\mathbf{0}$'s, and the \mathbf{u} 's are $T \times 1$ vectors of ones, zeros, and residuals, respectively; the π 's are scalars; and $\boldsymbol{\Pi}$ is a $k \times 1$ vector. Such an equation can be rewritten as

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{u}, \tag{A2}$$

where $\mathbf{Y}' = (\Delta \ln c', r^s, r^b)$ and \mathbf{u}' are $3T$ vectors, \mathbf{X} a $3T \times (k + 3)$ matrix, and $\boldsymbol{\beta}$ a $k + 3$ vector.

Given γ , an estimate of $\boldsymbol{\Pi}$ can be obtained by ordinary least squares from system (A1). One can then obtain an estimate of γ by grid search, minimizing the function $(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})'(\mathbf{X}'\mathbf{W}\mathbf{X})(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})$ for some given weighting matrix \mathbf{W} .

We adjusted this basic procedure to take into account the peculiarity of our problem and performed the following steps. (1) We estimate the reduced form (A1) for the three endogenous variables. (2) We compute the coefficient of the MA(1) process implied by the presence of sampling error in the residuals of the reduced form. (3) We forward filter equation (A2) and perform a grid search to get a first estimate of γ that minimizes $(\mathbf{Y}^* - \mathbf{X}^*\boldsymbol{\beta})'(\mathbf{X}'\mathbf{W}\mathbf{X})(\mathbf{Y}^* - \mathbf{X}^*\boldsymbol{\beta})$, where \mathbf{Y}^* and \mathbf{X}^* are the filtered \mathbf{Y} and \mathbf{X} , respectively, and \mathbf{W} is the identity matrix. (4) We obtain an estimate of the residuals and compute a new \mathbf{W} matrix based on these residuals: $\mathbf{W} = (\sum_i \mathbf{X}_i \mathbf{u}_i \mathbf{u}_i' \mathbf{X}_i')^{-1}$ and perform an additional grid search to determine our final estimate of γ . (5) The variance-covariance matrix of the parameters so obtained is computed using the formula

$$\mathbf{V} = \{[\sum (\mathbf{z}_i \otimes \mathbf{d}_i)][\sum (\mathbf{z}_i' \otimes e_i)(\mathbf{z}_i' \otimes e_i)']^{-1}[\sum (\mathbf{z}_i \otimes \mathbf{d}_i)]'\}^{-1},$$

where the e_i are the two final estimated residuals, and \mathbf{d}_i is the vector of derivatives of e_i with respect to γ , β_1 , and β_2 , that is,

$$\mathbf{d}_i = \begin{pmatrix} dc_i & dc_i \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

The Gauss program that computes these estimates is available from the authors on request.

References

- 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.
- Banks, James, and Johnson, Paul, eds. *How Reliable Is the Family Expenditure Survey? Trends in Income and Expenditure over Time*. London: Inst. Fiscal Studies, 1998.
- Blundell, Richard; Browning, Martin; and Meghir, Costas. "Consumer Demand and the Life-Cycle Allocation of Household Expenditures." *Rev. Econ. Studies* 61 (January 1994): 57–80.
- Browning, Martin; Deaton, Angus; and Irish, Margaret. "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle." *Econometrica* 53 (May 1985): 503–43.
- 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.
- Haliassos, Michael, and Bertaut, Carol C. "Why Do So Few Hold Stocks?" *Econ. J.* 105 (September 1995): 1110–29.
- Hansen, Lars Peter, and Singleton, Kenneth J. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectation Models." *Econometrica* 50 (September 1982): 1269–86.
- . "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *J.P.E.* 91 (April 1983): 249–65.
- . "Efficient Estimation of Linear Asset-Pricing Models with Moving Average Errors." *J. Bus. and Econ. Statis.* 14 (January 1996): 53–68.
- Kocherlakota, Narayana R. "The Equity Premium: It's Still a Puzzle." *J. Econ. Literature* 34 (March 1996): 42–71.
- Mankiw, N. Gregory, and Zeldes, Stephen P. "The Consumption of Stockholders and Nonstockholders." *J. Financial Econ.* 29 (March 1991): 97–112.
- Mehra, Rajnish, and Prescott, Edward C. "The Equity Premium: A Puzzle." *J. Monetary Econ.* 15 (March 1985): 145–61.
- Moffitt, Robert. "Identification and Estimation of Dynamic Models with a Time Series of Repeated Cross-Sections." *J. Econometrics* 59 (September 1993): 99–123.
- Vissing-Jørgensen, Annette. "Limited Stock Market Participation and the Equity Premium Puzzle." Manuscript. Chicago: Univ. Chicago, Dept. Econ., 1999.