

NBER WORKING PAPER SERIES

ASSET HOLDING AND
CONSUMPTION VOLATILITY

Orazio Attanasio
James Banks
Sarah Tanner

Working Paper 6567
<http://www.nber.org/papers/w6567>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 1998

This study has been financed under the ESRC Centre for the Microeconomic Analysis of Fiscal Policy at IFS. We are also grateful for funding from the EU under the Training and Mobility of Researchers scheme, contract no. ERBFMRXCT960016. Material from the FES made available by the ONS through the ESRC Data Archive has been used by permission of the controller of HMSO. Neither the ONS nor the ESRC Data Archive bear any responsibility for the analysis or interpretation of the data reported here. Thanks are due to Richard Blundell, Lucy Chennells, Lars Hansen, Costas Meghir, Guglielmo Weber, Frank Windmeijer for useful comments and discussions, as well as participants at the February 1997 TMR meeting on Savings and Pensions, the Summer 1997 NBER Workshop on consumption, ESEM 1997 and at seminars at Bristol University, IADB, INSEE, University of Warwick. The usual disclaimer applies. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1998 by Orazio Attanasio, James Banks and Sarah Tanner. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Asset Holding and Consumption Volatility
Orazio Attanasio, James Banks and Sarah Tanner
NBER Working Paper No. 6567
May 1998
JEL Nos. E21, G12

ABSTRACT

Recent studies have explored 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. We estimate ownership probabilities to separate "likely" shareholders from non-shareholders, enabling us to control for changing composition effects as well as selection into the group. We then construct estimates 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 non-shareholders, and more highly correlated with excess returns to shares. In particular, one cannot reject the predictions of the Consumption CAPM 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.

Orazio Attanasio
Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE
UNITED KINGDOM
and NBER
o.attanasio@ucl.ac.uk

James Banks
Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE
UNITED KINGDOM

Sarah Tanner
Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE
UNITED KINGDOM

1. Introduction

Asset pricing models based on the Euler equation for consumption have not performed well empirically. One of the reasons put forward has been the lack of variability of the intertemporal marginal rate of substitution of consumption, given plausible parameter values (see, for example, Hansen and Singleton, 1982, or Kotcherlakota, 1996, for a survey). Aggregate consumption growth, which in standard representative agent models 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 is 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) could be explained by the time series properties of consumption growth for asset market participants being systematically different from those of aggregate consumption growth.

This was a point stressed by Mankiw and Zeldes (1991). They find that a distinction between shareholders and non-shareholders is important for resolving the empirical failure of consumption-based CAPM models. But, the data they use contain information on food expenditure only. If food consumption is non-separable from the other components of consumption their estimates will be misleading.¹ Secondly, groups of shareholders and non-shareholders 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 shareowners *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 US, and particularly in the UK, where levels of share ownership and the composition of the group of shareholders have changed dramatically in recent years.

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 that allow us to identify groups of shareholders and non-shareholders over time, which also contain information on total consumption, do not exist. Instead we use a grouping estimator to repeated cross-section data and condition on *past* information to hold the composition of the group constant in looking at changes over time. This is an application of synthetic panel estimation which, to our knowledge, has not been used before. We define groups of shareholders and non-shareholders 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, i.e. we compute the IMRS between time t and $t+1$ using the consumption of households predicted to be shareholders at time t . This technique controls for changing composition of the group of shareholders between periods. It also solves the problem that the decision to own shares in each period is likely to be endogenous with respect to consumption.

The data are drawn from the UK Family Expenditure Survey, which has been collected continuously and consistently since 1968.² This gives us a long 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 twenty years also make the UK 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 flotation 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. These changes induced a change in the composition of shareholders and may also have changed the nature of the process generating prices and returns.³

¹ Even controlling for the non-separability of food from other items, the fact that food is a necessity does not make it ideal for testing the relationship between consumption growth and intertemporal prices, although this argument relates to power and efficiency, not consistency.

² In fact the data series we use in this paper starts in 1978 since that is when an education variable became available.

³ Allen and Gale (1994) construct a model of limited stock market participation and study its effect on asset prices volatility.

The plan of the paper is as follows. In the next section we review briefly the consumption CAPM model and its testable implications on the time series properties of asset prices. Section 3 presents evidence on recent changes in share-ownership in the UK and also shows asset returns over this period. Using the techniques developed in Hansen and Jagannathan (1991), we present mean-variance bounds on the IMRS computed from data on the returns to shares and Treasury Bills between 1978 and 1995. It can be seen clearly that the IMRS computed from aggregate consumption data over this period does not satisfy these bounds, generating an equity premium puzzle over this period. Section 4 discusses in detail the econometric technique we develop to characterise the time series properties of a variable for a group whose composition changes over time in a manner endogenous to the variable of interest. This technique has a number of potential applications. Section 5 compares the time series properties of consumption growth of ‘likely’ shareholders and non-shareholders as defined by the estimated probabilities and interprets the results. Section 6 concludes the paper with some remarks about potential extensions.

2. The Consumption CAPM model

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 (and omitting individual indexes for simplicity) the maximization problem can be written as follows:

$$\begin{aligned}
 & \max E_t \left[\sum_{s=t}^T U(C_s, v_s) \beta^{s-t} \right] \\
 (1) \quad & s.t. \\
 & \sum_{i=0}^n A_s^k = \sum_{i=0}^n A_{s-1}^k (1 + r_s^k) + Y_{s-1} - C_{s-1}
 \end{aligned}$$

where C_s denotes (non durable) consumption in period s , v denotes other factors that might affect the (marginal) utility of non durable consumption such as demographic variables, β is the discount factor, A^k is the amount of wealth held in asset k and r^k is the rate of return on that asset.

If asset k is held in periods t and $t+1$, a first order condition for this problem is

$$(2) \quad U(C_t, v_t) = E_t[\beta U(C_{t+1}, v_{t+1})(1+r_{t+1}^k)]$$

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:

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

where m_{t+1} is the intertemporal marginal rate of substitution between consumption in t and $t+1$, and n are the assets held by the consumer in non-zero 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. Note that equation (3) can also be expressed in terms of excess returns, for example between a risky and a safe asset:

$$(4) \quad E_t[m_{t+1}(r_{t+1}^i - r_{t+1}^j)] = 0 \quad j, i = 1, \dots, n$$

and will still formally identify the preference parameters.

A key implication of Consumption CAPM models, therefore, is that equilibrium returns are determined by a single factor: the IMRS. If there are complete markets and common information sets across all consumers, observed asset returns are sufficient to identify the IMRS. When either of these assumptions is violated one can derive relationships, implied by the observed asset returns, that impose restrictions on the unconditional moments of the IMRS. Hansen and Jagannathan (1991) derive bounds on the mean and variance of the benchmark portfolio from observed data on asset returns.

An alternative is to take a more structural approach based on the specification of a utility function. The orthogonality conditions implied by equation (3) or (4) for different assets can be used to estimate preference parameters in (1) and, provided the model is over-identified, test the over-identifying assumptions. This was the approach followed, for instance, by Hansen and Singleton (1982, 1983) who estimated several versions of equation (4) using aggregate time series data.

Equation (3) and (4) involve non-linear 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. Therefore, we log-linearize the Euler equation (3) under the assumption of CRRA (or iso-elastic) preferences to obtain:

$$(5) \quad \ln(1+r_{t+1}^j) = k_t^j + \gamma \ln\left(\frac{C_{t+1}}{C_t}\right) + \varepsilon_{t+1}^j, \quad j = 1, \dots, N$$

where ε_{t+1}^j is a term that includes expectational errors and changes in unobserved heterogeneity, γ is the inverse of the coefficient of relative risk aversion and k_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 log normality of consumption growth and asset returns, k_t^j is given by the following expression:

$$(6) \quad k_t^j = \ln \beta - \frac{\gamma^2}{2} \text{Var}_t[\Delta \ln C_{t+1}] - \text{Var}_t[\ln(1+r_{t+1}^j)] + \gamma \text{Cov}_t[\ln(1+r_{t+1}^j), \Delta \ln C_{t+1}]$$

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 ε_{t+1}^j and with the innovations to k_t^j , the parameters of equation (5) can be estimated using GMM techniques. In the absence of measurement error, and under the assumption of rational expectations, any variable dated $t-1$ or earlier is a valid instrument.

Considering equations (5) and (6) for two different assets, j and i , we can obtain a log-linear version of equation (4). After some manipulation, it is possible to derive:

$$(7) \quad E_t \left[\ln \left(\frac{1+r_{t+1}^j}{1+r_{t+1}^i} \right) \right] = \gamma \text{Cov}_t \left[\ln \left(\frac{1+r_{t+1}^j}{1+r_{t+1}^i} \right), \Delta \log C_{t+1} \right] - \text{Var}_t[\ln(1+r_{t+1}^j)] + \text{Var}_t[\ln(1+r_{t+1}^i)]$$

Notice that, since consumption growth enters only through its conditional covariance with asset returns, the identification of the curvature of the utility function is harder than in equation (5). If the conditional second moments on the right-hand-side of equation (7) are not predictable, neither are excess returns. Moreover, in such a situation one would not be able to identify the curvature parameter γ .⁴

Mankiw and Zeldes (1991) rely on an unconditional version of equation (7). Under the assumption of log-normality, taking unconditional expectations of equation (7) yields:⁵

⁴ Hansen and Singleton (1983) tested the predictability of excess returns and rejected the null very strongly. The advantage of working with excess returns is the fact that it does not require the measurement of inflation rates and marginal tax rates.

⁵ The terms in the variances of the returns are not included in equation (8) because on the left hand side we have the log of the mean rather than the mean of the log. Under the assumption of log-normality this allows us to eliminate the variance terms.

$$(8) \quad \log\left(\frac{E[1+r_{t+1}^j]}{E[1+r_{t+1}^i]}\right) = \gamma \text{Cov}[\tilde{r}_{t+1}, \Delta \log(C_{t+1})]$$

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

where $\tilde{r}_{t+1} = \log(1+r_{t+1}^j) - \log(1+r_{t+1}^i)$.


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 in equation (8). This approach is appealing because it relates the estimate of the parameter of interest directly to the time series properties (variance and correlation) of rates of return and consumption growth. However, it exploits a single orthogonality condition to estimate the unknown parameter and is therefore less efficient than the other methods discussed above. For this reason we do not use it.

3. Evidence on asset returns and asset ownership in the UK

3.1 Patterns of share ownership in the UK

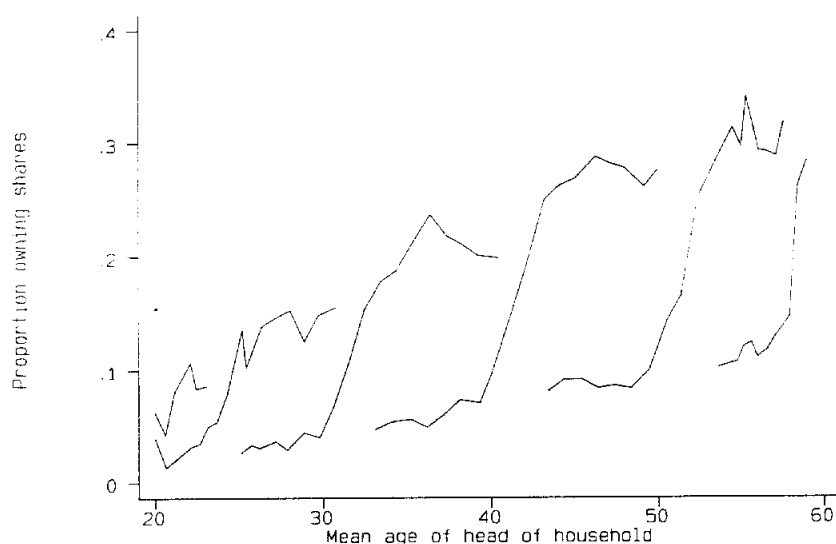
In keeping with the US, the overwhelming majority of UK households - more than 75 per cent - do not own shares directly. Given the substantial excess returns to shares (see Table 1), this is a puzzle (see Haliassos and Bertault, 1995, for a discussion of possible explanations). Furthermore, levels of share ownership greater than 20 per cent are a recent phenomenon.

Figure 1: Share ownership, 1978-1993



Levels of direct share ownership have changed dramatically in the UK over the past twenty years — more than trebling over the period 1985-1988.⁶ Recent patterns in share ownership are shown in Figure 1 using 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 (for a recent study see Banks and Tanner, 1996). These levels of ownership are very similar to those found in other data sources for individual years.⁷

Figure 2: Cohort profile – share ownership



The trends in aggregate ownership rates seen in Figure 1 mask significant differences in the experiences of different date-of-birth cohorts. In Figure 2 we present evidence on the share ownership of different cohorts. Each cohort's experience is plotted as a separate line, with the youngest appearing at the left hand side. The figure shows that the growth in share ownership happened throughout the age distribution (whilst being more pronounced for older households), implying that the life-time profiles for ownership for the younger cohorts have been shifted up at a much earlier age than their older counterparts.

⁶ It should be noted, however, that 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% and by 1994 it had fallen to 20.3%. This is largely the result of a rapid growth in institutional ownership by pension funds and insurance companies. In this paper the groups 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.

⁷ For example, the proportion owning shares in the 1988 General Household Survey was 21%.

The rapid increase in share ownership coincided with a number of measures designed by the Conservative Government 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). Privatisation accounts for a large part of the increase in the number of shareowners. 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 owned shares only in a privatized company. However, the evidence also suggests that there was a more general increase in ownership of shares. The information in the GHS shows that there was an increase in the proportion owning 'other' (i.e. non-privatised) shares from 10 per cent in 1987 to 13 per cent in 1988. Partly this may have been brought about indirectly through privatisation and the knock-on effects of increased awareness of and knowledge about share ownership.⁸ A second measure introduced by the Government was tax-favoured employee share schemes, three of which brought in between 1979 and 1984. By 1988 the total number of employees participating in such schemes was over 1.1 million.

Table 1
Share ownership by income and education

Year	Income quintile	<i>Highest level of education</i>		
		Compulsory schooling	A levels or equivalent	College education
1978	Quintile 1	2.0	8.1	11.3
	Quintile 2	2.2	6.9	8.5
	Quintile 3	2.7	11.0	11.4
	Quintile 4	3.9	6.5	13.0
	Quintile 5	6.2	19.6	29.6
1995	Quintile 1	7.6	15.7	18.9
	Quintile 2	13.9	25.8	22.8
	Quintile 3	20.1	27.7	31.5
	Quintile 4	22.4	37.1	31.1
	Quintile 5	28.9	40.5	42.4


⁸ 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 that this may not be effective in drawing stockholders from lower income groups.

Given the large changes in the level of ownership of shares and the differences across cohorts, it might be expected that the composition of the group of shareholders has changed over the period. 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 (as the results of the probit regression of share ownership in the next section will show more clearly). At the start of the period, for example, the average age of shareholders was 45. This had fallen to 43 by the end of the period. Similarly, the proportion of shareholders with no further education was 28 % at the start of the period, but had risen to 39% by the end. Over the same period, the proportion of the whole sample with compulsory education actually fell from 61% to 54%. Table 1 shows that the biggest increases in share-ownership came among those with high incomes, but low levels of education. Over the period there has been a considerable narrowing of the difference in ownership rates between education groups. In 1978, households with college education were more than four times more likely to own shares than those in the same income quintile with only compulsory education. By 1995, they were around twice as likely. If these changes in the composition of the group of shareholders imply differences in the consumption profiles of the group across the period, the implication of the consumption capital asset pricing market is that there will also be changes in asset market returns to preserve the equilibrium relationship. We explore this further below.

3.2 Asset returns

Real quarterly returns to shares and 3-month Treasury Bills are plotted in Figure 3 and summarized in Table 2. The share returns are given for the UK 500 share index. The main facts about the returns on shares and T-bills are not surprising. Over the period as a whole, as in the US, the share returns are substantially higher and more volatile than the returns on Treasury Bills. Given the large changes in share ownership that occurred during the mid 1980s, it is interesting to consider the behaviour of asset returns before and after these changes. The stock market crash of 1987 constitutes a natural breaking point both because it occurs at the end of the increase in ownership and because it allows us to define two periods that are not affected by a single large observation.

Figure 3: Quarterly real returns - shares and 3-month Treasury Bills



The excess return on shares is considerably lower in the second half of the period than in the first half. Also, share returns are slightly less volatile post 1987: both the standard error of real returns and that of excess returns are lower after the stock market crash. While the size of the two sample periods considered is too short to obtain any precise estimates of mean and variances of returns, other differences are also apparent. While the total variability of returns is lower after the increase in stock market participation, movements in returns and excess returns are less predictable. The R^2 of a simple OLS regression using lagged variables, such as share returns, T-bill returns, inflation rates and interest rate spreads is reduced from around 0.15 before 1987q1 to around 0.05 after 1987q4.

Table 2
Average quarterly returns

	Shares		T-Bills		'excess' return	
	Mean	<i>sd</i>	Mean	<i>Sd</i>	mean	<i>sd</i>
78q1-95q4	.0228	.0900	.0088	.0119	.0140	.0740
78q1-87q3	.0364	.0768	.0075	.0130	.0295	.0789
88q1-95q4	.0198	.0730	.0102	.0105	.0096	.0695


Following Hansen and Jagannathan, 1991, we use the observed data on asset returns to compute volatility bounds on the IMRS, expressed as mean-standard deviation pairs. This is given by the shaded area in Figure 4. For illustration, we also use quarterly data on

aggregate total expenditure from the UK National Accounts to compute estimates of the IMRS. Assuming that within-period utility functions exhibit Constant Relative Risk Aversion, and maintaining the assumption of inter-temporal separability, the IMRS is given by:

$$(9) \quad IMRS = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}$$

where γ is the coefficient of relative risk aversion and β the discount factor. Assuming different values of γ (between 0.5 and 5) and a discount rate of 2 per cent we plot the mean and standard deviation pairs for the IMRS implied by the growth in aggregate consumption for the period 1978-95. These are shown by the 'crosses' in Figure 4. This figure shows clearly that the IMRS mean-standard deviation pairs implied by aggregate expenditure data (and plausible values for the coefficient of relative risk aversion⁹ and of the discount factor) lie well outside the region admissible by the asset return data. In the next two sections we explore the possibility that limited participation in asset markets, and the stock market in particular, may be able to resolve this puzzle.

Figure 4
IMRS frontier computed using quarterly data



⁹ For example, Banks, Blundell and Tanner, 1996, obtain an estimate of the coefficient of relative risk aversion of -2.

4. Methods

4.1 The methodology

We want to look separately at the time series properties of consumption growth of shareholders and non-shareholders. But, since the FES is not a panel, we cannot 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. We need to use a grouping estimator to define groups of shareholders and non-shareholders over time.

A first obvious possibility, that is to group individuals *at each point in time* on the basis of current share ownership is likely to give misleading results because of changes in the composition of the group shareholders over time. Furthermore, share ownership is likely to be endogenous with respect to consumption. When hit by a shock, a household might decide to invest in stocks while it was not before (if the shock is positive) or might decide to liquidate its holding of stocks (if the shock is negative).

Our approach is to define the groups in terms of predicted ownership probabilities at a given point in time. Furthermore, 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). 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 non-shareholders 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 shareholders as:

$$(10) \quad \Delta \ln C'_{t+1} \equiv \left[\overline{\ln C_{t+1}} \mid \hat{p}(\text{Shareholder})_t > p_t \right] - \left[\overline{\ln C_t} \mid \hat{p}(\text{Shareholder})_t > p_t \right]$$

where $p(\text{Shareholder})_t$ is the predicted probability of owning shares and p_t is a cut off point. We compute a similar expression for the non-shareholders. We use the actual proportion of shareholders in our sample as the cut-off point in each time period. This implies that the cut-off point changes over the time period. But the groups defined in each subsequent time periods t and $t+1$ are formed on the basis of the same criterion: the probability of ownership at time t . 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 (10) for all pairs of adjacent time periods in our sample.

4.2 Problems

Three issues, arising from the complete lack of a longitudinal dimension in our data, need to be discussed. First, the theoretical model implies the consideration of share ownership at time t for individuals observed at time $t+1$, while we define the groups on the basis of *estimated probabilities of ownership*. Second, while the groups are defined consistently for any two subsequent periods, so that the definition of ‘consumption growth’ makes sense, group membership changes when considering different observations (over time) for the rate of consumption growth. Third, estimating the IMRS in equation (9) involves some aggregation problems that cannot be fully resolved given the lack of a longitudinal dimension. We discuss each of these issues in turn.

a) Estimated probabilities versus actual ownership

With access to panel data we would like to estimate:

$$(11) \quad \Delta \ln C_{t+1}^o \equiv \left[\overline{\ln C_{t+1}} | d_t = 1 \right] - \left[\overline{\ln C_t} | d_t = 1 \right] = \frac{E_i[\ln C_{t+1} d_t]}{E_i[d_t]} - \frac{E_i[\ln C_t d_t]}{E_i[d_t]}$$

where d is a dummy indicating share ownership and the subscript i indicates that the expectation is taken over the cross sectional dimension. But with repeated cross sections we are unable to compute the first of the two terms on the right hand side of equation (11). Having estimated a model for the probability of ownership, it might be tempting to weight individual log consumption at t and $t+1$ by the estimated probabilities. That is, one would approximate the right hand side of equation (11) by $E_i[\ln C_{t+1} E[d_t | z_t]] - E_i[\ln C_t E[d_t | z_t]]$, where z denotes the vector of variables used to model the probability of ownership. This procedure would only be appropriate if share ownership at t and log consumption at $t+1$, conditional on z_t , were uncorrelated, an assumption which is obviously not tenable.¹⁰

The expression for the consumption growth of likely shareholders on the right hand side of equation (10) differs from the expression on the right-hand side of equation (11) by four terms reflecting prediction errors from the probit. Two possible misclassifications can occur. Individuals who hold shares can have $\hat{p}(\cdot) < p_i$, and hence their consumption

is not counted in our measure of the IMRS. Similarly, individuals with $\hat{p}(\cdot) > p_t$, who do not hold shares will have their consumption falsely included.¹¹

$$\begin{aligned}
 \Delta \ln C_{t+1}^s &= \Delta \ln C_{t+1}^o + E[\ln C_{t+1} | \hat{p}(\cdot) > p_t, d_t = 0] \Pr\{d_t = 0 | \hat{p}(\cdot) > p_t\} \\
 &\quad - E[\ln C_t | \hat{p}(\cdot) > p_t, d_t = 0] \Pr\{d_t = 0 | \hat{p}(\cdot) > p_t\} \\
 &\quad - (E[\ln C_{t+1} | \hat{p}(\cdot) < p_t, d_t = 1] \Pr\{d_t = 1 | \hat{p}(\cdot) < p_t\} \\
 &\quad - E[\ln C_t | \hat{p}(\cdot) < p_t, d_t = 1] \Pr\{d_t = 1 | \hat{p}(\cdot) < p_t\})
 \end{aligned}
 \tag{12}$$

If the Probit discussed above predicted share ownership perfectly, $\Delta \ln C_{t+1}^s = \Delta \ln C_{t+1}^o$ and we would be measuring the consumption growth of actual shareholders *at time t*. When this is not the case, one wants to establish whether the fact that the last four terms in (12) are not zero introduces a bias in the estimates of the structural parameters and in the tests of over-identifying restrictions. To do so, one has to evaluate the covariance between these (unobserved) terms and the instruments used in the GMM procedure described below. As we use instruments that are lagged two periods our procedure should still yield consistent estimates unless there are reasons to believe that the terms on the right hand side of (12) exhibit serial correlation.

Three points should be stressed. First, even when panel data are available, if ownership is not perfectly predictable because of the presence of unobserved heterogeneity, using the consumption growth of actual owners might introduce important biases, as we discuss below. Second, using the consumption of *actual* share holders at time t and that of *predicted* share holders at time $t+1$ would be inappropriate as the two groups would not be homogenous. Our procedure, instead, defines groups consistently between adjacent time periods. Finally, our procedure is, at the very least, a test of the null hypothesis that limited stock ownership is not the explanation of the empirical failure of the consumption CAPM. We check whether the consumption behaviour of what we define as the likely shareholders is systematically different from the rest of the population.

b) Composition effects

The second issue is that the composition of the group of likely shareholders changes over time, both because the probability of ownership may change over time and because

¹⁰ Such an assumption is particularly unappealing when the ζ 's are as parsimonious as those we use.

the cut-off point changes. If intertemporal prices were the same across individuals and the utility function depended only on consumption and not on unobserved (or unaccounted for) heterogeneity, this would not be a problem. The measured IMRS would differ from the expected one only because of an expectational error that would average to zero over time. In the presence of unobserved heterogeneity, however, the measured IMRS encompasses both genuine changes in consumption growth and composition effects. This introduces a spurious source of volatility.¹² To make this point clear, suppose that the instantaneous utility function is given by $U(C_t^h, v_t^h) = (1-\gamma)^{-1} (C_t^h)^{1-\gamma} \exp(v_t^h)$, where v_t^h represents unobserved heterogeneity and we are ignoring the effect of demographic and other observable variables for notational simplicity. The IMRS corresponding to this utility function is: $IMRS = (C_{t+1}^h)^{-\gamma} (C_t^h)^\gamma \exp(v_t^h - v_{t-1}^h)$. If v_t^h is constant over time and the groups are formed consistently, as our procedure requires, the presence of unobserved heterogeneity does not create any problems, as the terms in v_t^h would drop out of the IMRS. On the other hand, if v_t^h is a random walk with a group specific drift, changes in the compositions of the groups might create some serious distortion in the characterization of the time series properties of the IMRS. For this reason, it might be important to check whether the results change dramatically when we change the way we select the ‘cutoff points’ in our procedure.

As we mentioned above, it should be noted that the presence of unobserved heterogeneity in the utility function is likely to induce serious problems even when panel data are available, if one uses actual ownership. When the unobserved taste shocks v_t^h are correlated with the decision to hold stocks (as is likely), then using individual data and individual specific instruments (even if lagged a few periods) will produce inconsistent estimates. Like any grouping estimator our procedure has an Instrumental Variable interpretation. Averaging over the individuals belonging to a group (defined by a non-linear function of deterministic variables) avoids the biases caused by individual specific fixed effects. It should be stressed that our procedure relies only on temporal variability

¹¹ If one believed that selection into the group was exogenous to consumption growth the two last terms in equation (10) could be eliminated by using actual, rather than predicted, data in period t .

¹² While we correct for small sampling error, we are now referring to systematic changes in the group of shareowners.

to identify the parameters of interest. If there are group specific fixed effects, they can be absorbed in the constant of the equation.¹³

c) Aggregation issues

The Euler equations used in asset pricing relationships are not necessarily a simple function of consumption growth. Typically, they involve the rate of growth of the marginal utility of consumption. As long as the IMRS can be made a linear function of parameters, this is not necessarily a problem. Indeed, the Euler equations we estimate below (and discuss in Section 2) were log-linearized for this reason. However, non-linearities do constitute a problem if one works with expressions for the IMRS such as that in equation (9) used to compute the Hansen-Jagannathan bounds. This expression involves the interaction of consumption at time t and $t+1$, which we cannot compute. Therefore, we approximate the left-hand side of (9) by the following expression:

$$(13) \quad IMRS = \beta \frac{\overline{(C_{t+1})^{-\gamma}}}{\overline{(C_t)^{-\gamma}}}$$

Taking the ratio of the averages, rather than the average of the ratios obviously introduces some biases. Unfortunately, there is not much we can do about this issue, which may introduce biases of unknown nature. However, we can derive conditions under which there is no bias in following this approximation. If, for instance, the rate of growth of individual consumption is uncorrelated with the initial level of consumption, and consumption is log-normally distributed in the cross section, it can be showed that our approximation does not introduce any bias.¹⁴

5. Results

5.1 A Probit model for share ownership

We obtain the probabilities of share ownership by estimating a probit model on a pooled sample of data containing more than 80,000 households. On the right hand side we include polynomials in age and time, education dummies and interaction terms in these variables. It is important to stress that the time trends are interacted with the other

¹³ This is important because, given the structure of the model, the instruments are only weakly exogenous. Another problem that can be dealt with by our grouping procedure is that of possible mis-classification of share ownership. This can arise because ownership is inferred from dividend income in the previous 12 months.

¹⁴ Attanasio and Weber (1993) and Constantinides and Duffie (1996) discuss similar issues, in a different context. Constantinides and Duffie (1996) use the same assumption to derive some of their results. Another assumption which

explanatory variables, to allow for the fact that the effects of factors such as age and education appear to change over time. The results (reported in Table 3) 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 diminishes over time. We have obtained very similar results by estimating a different probit for each year in the sample.

Table 3
Results of probit estimation

	<i>Parameter</i>	<i>Standard error</i>	<i>Marginal effect</i>
Age of head	0.281	0.022	0.055
(Age of head) ²	-0.021	0.006	-0.004
Head has A levels	0.697	0.031	0.160
College education	1.198	0.039	0.349
Age*Alevels	0.082	0.013	0.016
Age*College	0.117	0.016	0.023
Trend	3.522	0.875	0.687
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	-0.212	0.083	-0.042
Age*Trend ²	-0.365	0.047	-0.071
Alevels*Trend	-0.634	0.058	-0.124
College*Trend	0.122	0.076	0.024
Constant	-2.022	0.051	—
N	83,736		
Pseudo-R ²	0.1203		

5.2 Results on 'predicted shareholders'

Defining groups of shareholders and non-shareholders on the basis of their predicted probabilities of share ownership, we compare the time series properties of their consumption growth rates.¹⁵ We remove both durable spending and housing costs from

would guarantee the absence of aggregation biases is that higher moments of the cross sectional distribution of consumption are constant over time.

¹⁵ Because the size of each cell used to compute average consumption is not very large, sampling error induces part of the time series variability of consumption growth. Moreover, as the group of 'predicted share holders' is substantially smaller than the group of 'unlikely' share holders, it is important to control for this problem. Fortunately, one can use the information on within cell variances to correct the estimate of the total variance, as in the following expression:

$$Var(\Delta \bar{x}_i^c) = Var(E\Delta x_i^c) + \left[\frac{1}{N_i^c} Var_i(x_{it}^c) + \frac{1}{N_{i-1}^c} Var_i(x_{it-1}^c) \right]$$

the measure of total weekly expenditure in the FES. Since we look at the growth of consumption between quarters, we take out seasonal effects and also adjust for household size using equivalence scales estimated from FES data (see Banks and Johnson, 1994). On average, the predicted shareholders have higher consumption growth over the period than non-shareowners — 0.26 per cent per quarter for shareholders compared to 0.04 for non-shareholders. The standard deviation of the time-series of average consumption growth (adjusted to take account of the variance of within-cell measurement error) is one and half times as high — 3.7 for shareholders compared to 2.5 for non-shareholders.¹⁶ This suggests that differences between shareowners and non-shareowners are likely to be important in resolving the equity premium puzzle. But, resolving the puzzle requires not only the variance of consumption growth of shareholders to be larger than that of non-shareholders, but also a higher correlation of the IMRS of shareholders with the excess returns on shares. More generally, one can use any of the techniques discussed in section 2.

There are several ways in which the IMRS can be estimated or approximated. In the simplest version of the life cycle model, the marginal utility of consumption is a monotonic transformation of consumption. In more complex and possibly realistic versions of the model, the IMRS might depend in a flexible fashion on the composition of the household as well as on labor supply behavior. As a detailed characterization of preferences is not the main focus of this paper, we use a parsimonious specification. We assume that utility is an isoelastic function of non-durable consumption per adult equivalent.


We use the time series of consumption growth of the predicted group of shareholders to construct mean-variance pairs for the IMRS, assuming values for the coefficient of relative risk aversion between 0.5 and 5 and a discount rate of 2 per cent. These are shown by the ‘crosses’ in Figure 5a. Clearly — in comparison to using aggregate expenditure data — the plausible estimates of the IMRS mean-standard deviation pairs implied by the time series of consumption growth of shareholders lie much closer to the

where $E\Delta x_t^c$ is the population mean whose variance we are interested in studying, Var_t is the cross sectional variance and N_{ct} is the cell size at time t for group c .

¹⁶ If we were to compute the rate of growth for actual shareowners and non-shareowners, the difference is even more dramatic. The standard deviation of the time-series of average consumption growth (adjusted to take account of the


IMRS frontier implied by the observed asset returns. It is also the case that, given the smaller variance of consumption growth for the group of likely non-shareholders, we need a much larger value of γ to get close to the HJ bounds for this group (see Figure 5b). These results suggest that distinguishing between shareholders and non-shareholders is crucial to resolving the equity premium puzzle.

Figure 5a
Predicted shareholders



variance of within-cell measurement error) is almost twice as high — 6.52 for shareholders compared to 3.37 for non-shareholders.

Figure 5b
Predicted non-shareholders



Our second approach is to estimate versions of the Euler equation given by equation (5) for the real return on shares for the total sample and for the two groups defined on the basis of the conditional probability of share ownership. Results are reported in Table 4. As mentioned in Section 2, in the absence of measurement error, any variable dated $t-1$ is a valid instrument. As we use grouped data, however, small sample variability induces MA(1) errors and therefore one has to lag consumption growth twice to avoid getting inconsistent estimates.¹⁷ Furthermore, one needs to correct the standard errors for the presence of such an error structure. Finally, one can improve the efficiency of the estimates relative to a simple IV procedure by using a GMM estimator.

Because the estimated variance covariance matrix of the residuals after the first step is not always positive definite,¹⁸ we cannot always use a standard GMM estimator. Therefore, we develop a procedure, which is slightly different from the standard one. The estimator we use, whose details are given in the Appendix, consists of three steps. A

¹⁷ The presence of time aggregation effects also induces MA(1) residuals.

¹⁸ As discussed in the Appendix, because of the nature of our residual, the use of a Newey-West estimator to guarantee that the variance covariance matrix is positive definite is not legitimate.

first step that gives consistent estimates; a second that corrects for the presence of MA residuals and a third that computes an efficient GMM estimation on the transformed model.

As before, consumption growth is measured as the log change in de-seasonalised 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 the table. The only instrument that deserves a mention is a financial market liberalization dummy, which is meant to capture the process of transformation of British financial markets in the second half of the 1980s. Results are not greatly affected by the use of this particular instrument.

The results for the group of likely shareholders, reported in the first column of Table 4, are remarkably good. The point estimate of the coefficient of the iso-elastic utility function is consistent with concavity. It implies a value of the elasticity of intertemporal substitution of about 0.80, consistent with other estimates of this parameters in the UK reported in the literature.¹⁹ Finally, the value of the test of over-identifying restrictions is very low and does not indicate any deviations from the null.

Table 4
GMM estimation of Euler equation for the return on shares

	Likely Shareholders	Unlikely Shareholders	Whole sample
Constant	0.0165 (0.012)	0.040 (0.011)	0.034 (0.009)
γ	1.238 (0.513)	-2.694 (1.449)	0.0331 (1.157)
Sargan test (p-value)	1.456 (0.962)	7.484 (0.279)	12.19 (0.059)

Note: Standard errors in parentheses. Instruments include a constant, a dummy for financial liberalization and the second lag of: the rate of growth in consumption, share returns, t-bills return, the spread between 3-month t-bills and 20 year bonds, the spread between 20 years and 5 years bonds and inflation. The sample period is from 1978q4 to 1995q4. The Sargan test is a test of the over-identifying restrictions. Details on the technique implemented to obtain these estimates are in the Appendix.

The point estimates of γ obtained for the group of unlikely shareholders and for the total sample are shown in the second and third columns of Table 4. They are either very small (for the whole sample) or negative (for the unlikely shareholders). For the whole sample, the test of over-identifying restrictions has a p-value just above 0.05.

In the empirical asset pricing literature, researchers often reject the over-identifying restrictions implied by the Euler equation for consumption when they consider simultaneously more than one asset. These results are the counterpart to the fact that the first two moments of the IMRS based on aggregate consumption are outside the Hansen Jagannathan bounds. In Table 5 we report the results obtained estimating simultaneously the Euler equation for shares and T-bills returns.

Table 5
GMM estimation of Euler equation for the return on shares and T-Bills

	‘Likely’ Share owners	‘Unlikely’ share owners	Whole sample
γ	0.3936 (0.181)	-0.4728 (0.331)	0.0492 (0.140)
Const. (Shares)	0.0372 (0.010)	0.0311 (0.009)	0.0378 (0.009)
Const. (T-Bills)	0.0108 (0.004)	0.0182 (0.003)	0.0094 (0.0013)
Sargan test (p-value)	11.7 (0.556)	24.3 (0.029)	24.7 (0.025)

Notes: See Table 4.

As in Table 4 the results for the group of likely shareowners are very different from those of the other two groups and much more consistent with the predictions of the theory. While the coefficient of the iso-elastic utility function is lower than in the previous case, it is positive and significantly different from zero, even though is not estimated with a great precision. Furthermore, the test of over-identifying restrictions never rejects the null. The message that emerges from the first column of the Table is that even one considers two assets simultaneously for the group of likely shareholders,

¹⁹ See, for instance, Attanasio and Weber (1993), Banks, Blundell and Preston (1994) and Blundell, Browning and Meghir

one does not reject the over-identifying restrictions implied by the theory. Furthermore, one obtains point estimates of the parameter of interest that are not inconsistent with a concave utility function.

The evidence for the other two groups considered is, as in Table 4, quite different. The lack of precision of our estimates does not mean, for these two groups, that our procedure is unable to reject the orthogonality restrictions implied by the theory. On the contrary, for both groups we obtain strong rejections of the over-identifying restrictions. For the group of ‘unlikely share holders’, the estimated coefficient of relative risk aversion is again negative, even though is not estimated with any precision.²⁰

6. 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 UK. Since the first order conditions for the model only hold as an equality for individuals that are currently participating in asset markets, it is natural to look at the consumption behaviour 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 also we 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 data that are available, 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 present an extension of the synthetic cohort technique which defines groups of individuals with constant membership at adjacent dates on the basis of the estimated probabilities of owning stocks.

(1994) who estimate an Euler equation for consumption using returns on Building Societies deposits

²⁰ Estimating the Euler equation for the rate of returns on T-Bills only yields parameter estimates measured with very little precision for all groups. These results are somewhat at variance with the evidence presented in the UK literature estimating Euler equation with synthetic cohort data and relatively safe rates of returns. (see, for instance, Attanasio and Weber, 1993, Banks, Blundell and Preston, 1994, Blundell Browning and Meghir, 1994). The two main differences that are likely to explain the lack of precision are the different sample period and the use of T-Bill rates instead of the rate on Building Societies deposits typically used in that literature. Changes in the late 1980s in the published series prevented us from using that rate. The sample period in the papers mentioned above included the early 1970s when real rates of return on ‘safe’ assets exhibited a substantial amount of variability.

We obtain strong results. Firstly, the first two moments of the Intertemporal Marginal Rate of Substitution for the group of likely shareholders are remarkably close to the Hansen-Jagannathan bounds derived from the time series properties of returns on shares and Treasury Bills, which is not the case with aggregate data. Second, when we estimate Euler equations for the same group, both using a single asset and two assets simultaneously, we obtain sensible values for the parameter of interest (the coefficient of relative risk aversion) and we fail to reject the over-identifying restrictions implied by the model. Finally, for the other groups (the total sample and the unlikely shareholders, we either obtain unappealing estimates of the structural parameters (violating concavity of the utility function) or rejections of the over-identifying restrictions. This last result is important in showing that there is some empirical power in our approach. To summarise, we have shown that the time series properties of the consumption of shareholders are very different from those of aggregate consumption. And they are different in a way that is consistent with the implications of the Consumption CAPM.

The results suggest a number of extensions. Given the limitations in our sample period and the reliance on T-asymptotics to identify the parameters of interest, we have worked with very simple preference specifications. It would be interesting to work with preferences that are a more general the way demographic and labour supply factors are allowed to affect utilities. One possibility would be to estimate these effects using a longer time period and the Euler equation for a relatively safe asset and then check over the shorter period whether the orthogonality conditions hold, given the particular preference structure estimated. 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. The problem with the latter, however, is that they are very hard to study without longitudinal data.

The most important challenge, 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 privatisations of the mid 1980s, and by the associated publicity, the trends cannot be explained only by that episode and/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.

7. References

- Allen, F. and D. Gale, 1994, 'A Model of Limited Stock Market Participation', *American Economic Review*, 84, 933-55.
- Attanasio, O., 1991, 'Risk, Time Varying Second Moments and Market Efficiency', *Review of Economic Studies*, 58, .
- Attanasio, O. and M. Browning, 1995, 'Consumption over the Business Cycle and over the Life Cycle', *American Economic Review*, 85, 118-1127.
- Attanasio, O. and G. Weber, 1993, 'Consumption growth, the interest rate and aggregation', *Review of Economic Studies*, 60(3), pp. 631-649.
- Banks, J. and P. Johnson, 1994, 'Equivalence scale relativities revisited', *Economic Journal*, 104(425), pp. 883-890.
- Banks, J., R. Blundell and I. Preston, 1994, 'Life-Cycle Expenditure Allocations and the Consumption Cost of Children' *European Economic Review*, 38, 1234-78.
- Banks, J., R. Blundell and S. Tanner, 1996, 'Is there a retirement savings puzzle?' UCL discussion paper 96-25, forthcoming in *American Economic Review*
- Banks, J. and S. Tanner, 1996, 'Savings and wealth in the UK: evidence from micro data', *Fiscal Studies*, 17(2), 37-64.
- Blundell, R., Browning, M. and C. Meghir, 1994, 'Consumer Demand and the Lifetime Allocation of Consumption,' *Review of Economics Studies*, 61, 57-80.
- Browning, M., Deaton, A. and M. Irish, 1985, 'A Profitable Approach to Labor Supply and Commodity Demand over the Life Cycle', *Econometrica*, 53, 503-44.
- Constantinides G.M. and D. Duffie, 1996, 'Asset Pricing with Heterogeneous Consumers', *Journal of Political Economy*, 104, 219-40.
- Deaton, A., 1985, 'Panel Data from Time Series of Cross Sections,' *Journal of Econometrics*, 30, 109-26.
- Epstein, L. and S. Zin, 1989, 'Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns. A Theoretical Framework', *Econometrica*, 57, 937-69.
- Halliassos, M., and C.C. Bertaut, 1995, 'Why do so few people hold stocks?', *Economic Journal*, vol. 105, no. 432, 1110-1129.
- Hansen, L.P., 1982, 'Large Sample Properties of Generalized Method of Moments Estimations in Models with Rational Expectations', *Econometrica*, 50, 1029-54.

- Hansen, L.P. and R. Jagannathan, 1991, 'Implications of security market data for models of dynamic economies', *Journal of Political Economy*, 99(2), 225-62.
- Hansen, L.P. and F. Richard, 1987, 'The Role of Conditioning Information in deducing Testable Restrictions Implied by Dynamic Asset Pricing Models', *Econometrica*, 55, 587-613.
- Hansen, L.P. and K.J. Singleton, 1982, 'Generalized instrumental variables estimation of nonlinear rational expectation models,' *Econometrica*, 50, 1269-86.
- Hagiwara, M. and M. Herce, 1997, 'Risk Aversion and Stock price sensitivity to dividends', *American Economic Review*, Vol 87, No. 4 (November), 738-745.
- Hayashi, F. and C. Sims, 1983, 'Nearly Efficient Estimation of Time Series Models with Predetermined, But Not Exogenous, Instruments', *Econometrica*, 51, pxx.
- Heaton, J. and D. Lucas, 1996, 'Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing', *Journal of Political Economy*, 104, 443-87.
- Kotcherlakota, N.R., 1996, 'The equity premium: It's still a puzzle', *Journal of Economic Literature*, 34(1), 42-71.
- Mankiw, N.G. and S. Zeldes, 1991, 'The consumption of stockholders and non-stockholders' *Journal of Financial Economics*, 29, 97-112.
- Mehra, R., and E.C.Prescott, 1985, 'The equity premium: a puzzle', *Journal of Monetary Economics*, 15(2) 145-61.
- Moffitt, R., 1993, 'Identification and Estimation of Dynamic Models with a Time Series of Repeated Cross-Sections', *Journal of Econometrics*, 59, 99-124.
- Newey, W. and K. West, 1987, 'A Simple Positive Semi-Definite Heteroscedasticity and Autocorrelation Consistent Covariance Matrix', *Econometrica*, 55, 703-8.
- Poterba, J. and A. Samwick, 1996, 'Stock ownership patterns, stock market fluctuations and consumption', NBER reprint No. 2027.

Appendix: The GMM procedure used in Tables 4 and 5.

In this appendix we discuss the GMM technique we have used to obtain the results reported in Tables 4 and 5. Since Hansen (1982) and Hansen and Singleton (1982) it has become customary to use GMM estimators to exploit the orthogonality conditions implied by asset pricing relationships to obtain estimates of the structural parameters of the model and test its validity. Following the notation in Hansen and Singleton, our model can be represented by the following set of nk orthogonality conditions:

$$(A1) \quad E[u_t \otimes z_t] = 0$$

where z_t is a vector of instruments and u_t a vector of residuals of dimension equal to the number of assets for which we are considering the Euler equations. A consistent estimator is obtained by minimizing the following expression.

$$(A2) \quad g_T' W g_T \quad \text{where} \quad g_T = \frac{1}{T} \sum_t u_t \otimes z_t$$

and W is a positive definite weight matrix. Efficiency in the use of the GMM is achieved by the use of a weighting matrix that reflects both the presence of heteroscedasticity and autocorrelation. Since the researcher is often ignorant about the form of autocorrelation, and or heteroscedasticity, it is desirable to use an estimate of such matrix robust to the presence of a wide variety of time varying second moments. In particular, the residuals of a first step estimation might be used to construct the following matrix:

$$(A3) \quad S = P_0 + P_1 = \sum_t g_t g_t' + \sum_j \sum_t g_t g_{t-j}' + \sum_j \sum_t g_{t-j} g_t'$$

where the sums in j reflect the presence of autocorrelation and run, theoretically, up to infinity. In practice, the sums in j run to some fixed number and, in small samples, there is no guarantee that the matrix S is positive definite. Newey and West (1987) derive a correction to the A3 which down-weighting the elements of P_1 by triangular weights in a way that guarantees positive definiteness. Essentially, their estimator is consistent if one assumes that the triangular weights become asymptotically rectangular.

In our application, however, we have a substantial amount of information on the nature of the residuals of our Euler equations. Under the null, the residuals are made of two components, a white noise (or MA(1) if time aggregation is an issue) component reflecting expectational errors and a MA(1) component with a negative unit root reflecting measurement error and small sample variability deriving from the use of cohort

grouped data. The overall residual is therefore likely to be an MA(1) with a coefficient whose size depends on the relative variance of its components.

In such a situation the sums in j in the equation above should collapse to a single term and therefore the use of the Newey-West estimator is not justified. In particular, imposing an arbitrary weight on P_1 to make the whole matrix positive definite would imply that the estimates we obtain are not necessarily efficient and that our asymptotic inference would be biased.

A preliminary analysis of the estimated (first step) residuals, confirms that they are indeed well represented by an MA(1) process. Furthermore, looking separately at the two equations and at the various groups yields results that are consistent with our interpretation.²¹ While knowing the autocorrelation structure of the residuals is obviously an advantage, it leaves us with the possibility of obtaining matrices that are not positive definite in small samples. Such a problem does occur especially when we consider the two assets jointly. We therefore derive an alternative procedure that trades off some of the generality (A3) above with a simpler structure that is more likely to yield positive definite estimates of the variance covariance matrix (and therefore of the weighting matrix used in the second step of the GMM estimator).

Our procedure consists of three steps. In the first, we get consistent estimates of the parameters by minimizing the expression in (A2) with W given by the inverse of $Z'Z$ (or in the case of two assets) a block diagonal matrix with the inverse of $Z'Z$ on the diagonal. We then use the estimated residuals to obtain an estimate of the first order autocorrelation coefficients for the residuals of our model (we have one coefficient for each asset). With these coefficients we construct TxT matrices with ones on the diagonal and the first order autocorrelation coefficient in the $[k,k-1]$ and $[k,k+1]$ elements. We then compute the upper triangular Choleski decomposition of the inverse these matrices and pre-multiply the variables (but not the instruments) in our equations by such matrices. This is equivalent to an application of the techniques suggested by Hayashi and Sims (1983) and transforms the original model into one with no autocorrelation in the

²¹ For instance, we find that the MA coefficient of the share return equation is smaller in size, consistent with the fact that the variance of the expectational error for this equation is larger than the one of the T-Bills equation. Furthermore, the size of the MA coefficient decline with the size of the cells.

residuals. It is important that the instruments are not filtered and that the matrix used to filter the variables in the model is upper, not lower, triangular.

The second step of the procedure estimates the transformed model using the same W matrix as in the first step. The standard errors of the estimates obtained in this step can be estimated by the following expression:

$$(A4) \quad V = (DWD')^{-1} DWP_0WD'(DWD')^{-1} \text{ where } D = \frac{\partial g_T}{\partial \theta}$$

and θ is the vector of estimated parameters. In the final step, we re-estimate P_0 , and use its inverse to construct a new weight matrix to be used in expression A2. In other words, we first filter the model to eliminate the presence of autocorrelation. We then apply GMM (allowing for the presence of heteroscedasticity of unknown form) to the transformed model under the assumption that there is no autocorrelation.

Two comments on the application of this technique are in order. Firstly, in the cases in which the coefficients are reasonably well determined, the point estimates do not vary dramatically in the various steps. Second, when the standard GMM procedure can be applied, in that the estimated S matrix is positive definite, the results did not differ substantially from the ones we obtain with our procedure.