Consumption, Wealth and Expected Stock Returns in Australia*

Lance A. Fisher**
School of Economics
University of New South Wales
Australia

Graham M. Voss
Department of Economics
University of Victoria
Canada

Preliminary

Abstract

Following Lettau and Ludvigson (2001), this paper shows that a measure of private dissaving which accounts for capital gains can potentially predict real stock returns and growth in consumption and non-asset income. Private dissaving is measured as the deviation from the long-run cointegrating relationship between the logs of consumption, non-asset income and wealth. Empirically, private dissaving is found to predict real stock returns and risk premia on stocks in Australia over short and intermediate horizons. It is not found to predict consumption growth at any forecast horizon, consistent with the permanent income hypothesis.

Key words: consumption, wealth, asset returns, permanent income, VAR

JEL Classification: C32, E21

*Corresponding author:
Lance A. Fisher
School of Economics
University of New South Wales
Sydney, NSW, 2052
Australia
E-mail: L.Fisher@unsw.edu.au

**We would like to thank Brett Winton and Ellis Connolly of the Reserve Bank of Australia for assistance with the data used in this study and Geoffrey Kingston and Paul Richardson for helpful comments.

Consumption, Wealth and Expected Stock Returns in Australia

1. Introduction

The interaction between macroeconomic variables and asset returns is an important focus of study in both macroeconomics and finance. In an important paper, Lettau and Ludvigson (2001) find that a measure of private dissavings that takes account of capital gains is very successful in forecasting both stock returns and excess stock returns over short and intermediate horizons in U.S. data. Prior to this study, only moderate success had been achieved in forecasting stock returns. That limited success was achieved using financial market variables, for example, term premiums and short-term interest rates (Campbell (1987)) while real macroeconomic variables were not found to be particularly useful in forecasting returns[1].

The success of private dissaving as a predictive variable for stock returns derives from a consideration of how forward-looking agents behave. Lettau and Ludvigson (2001) show that for forward-looking agents, private dissaving reflects their expectations of future real asset returns and consumption growth. Under the permanent income hypothesis, consumption growth is not expected to change much so private dissaving signals primarily agents' expectations of future real asset returns. Lettau and Ludvigson (2001) define private dissavings as the deviation from the long-run or cointegrating relationship between consumption, non-asset income and wealth. When asset returns are expected to be higher in the future, forward-looking agents respond by increasing consumption out of current non-asset income and wealth, causing private dissaving to increase. Higher future returns allow agents' to increase current consumption while maintaining consumption levels expected to prevail in the future. Thus, an increase in

private dissavings signals expectations of higher future asset returns. In the same way, a decrease in private dissaving signals expectations of lower future asset returns.

In this paper, we derive an expression relating private dissaving to expected future real asset returns, consumption growth and non-asset income growth. Growth in non-asset income appears explicitly in this expression because we do not define aggregate wealth to include the market value of tradeable human capital as done by Lettau and Ludvigson (2001). In particular, we assume human capital is non-tradeable. The advantage of this formulation is that it can be related to early tests of the permanent income hypothesis (Campbell (1987)). In this formulation, private dissaving can potentially signal movements in expected future asset returns and consumption and non-asset income growth. As in Lettau and Ludvigson (2001), private dissaving can be thought of as the deviation from a long-run consumption function that relates consumption to non-asset income and wealth. For Australia, Tan and Voss (2003) found evidence for a long-run or cointegrating relationship between per-capita levels of non-durables consumption, non-asset income and wealth for the sample 1988(4)-1999(3). This paper builds on their work and finds evidence for a cointegrating relationship between the log-levels of per-capita non-durables consumption, non-asset income and wealth for the extended sample 1988(4)-2002(1). We depart from Tan and Voss and use log-levels since this is how the variables appear in the expression that forms the basis for our empirical analysis. Private dissaving is defined as the deviation from the estimated cointegrating relation between the logs of consumption, non-asset income and wealth. Throughout the paper, it is denoted $\hat{cay}_t$ (for time t) as in Lettau and Ludvigson (2001).

The expression we derive relating private dissaving to expected future real asset returns and consumption and non-asset income growth forms the basis for the empirical analysis. This analysis investigates the predictive content of $c\hat{a}y_t$ for future asset returns and for future growth in consumption and non-asset income. We find that $c\hat{a}y_t$ has predictive power for both real stock returns and excess returns on the Australian ASX200 accumulation index and that this predictive power is strongest for returns over horizons of up to a year. Our results show that the risk premium on stocks is time varying over short and intermediate horizons. We also find that $c\hat{a}y_t$ has no predictive power for growth in consumption or non-asset income at any forecast horizon.

Brennan and Xia (2002) take issue with Lettau and Ludvigson's framework in which the predictive power of $c\hat{a}y_t$ for stock returns derives from the forward-looking behaviour of agents who take account of future investment opportunities in planning their consumption streams. They argue that the predictive power of $c\hat{a}y_t$ is mainly due to "look-ahead bias" because the coefficients in the cointegrating relation are estimated from the full sample. To emphasize this point, they show that when consumption is replaced with calendar time in the cointegrating relation, the resulting variable "$t\hat{a}y_t$" forecasts stock returns just as well. Since *t* is just a linear time trend and has no foresight their argument goes, the predictive power of $t\hat{a}y_t$ must be attributable to look-ahead bias.

Lettau and Ludvigson (2002) present a persuasive reply to this critique. They argue that to test the implications of a framework where consumers are forward-looking, the appropriate estimation strategy is to use the full-sample estimates of the coefficients in the cointegrating relation between consumption, non-asset income and wealth. This is

because only the full-sample estimates are closest to the true parameters that would have been known to the representative consumer when making consumption and investment decisions. The theoretical framework shows that $c\hat{a}y_t$ has predictive power for asset returns provided it accurately reveals the deviation from the common stochastic trend in consumption, non-asset income and wealth. That requires accurate estimates of the coefficients in the cointegrating relation, which in turn requires a long sample of data. Furthermore, they argue that it is not surprising that $t\hat{a}y_t$ forecasts stock returns as it can be seen as a proxy for $c\hat{a}y_t$ because consumption is a very smooth series. In fact, they show that when consumers are forward looking, desire to smooth consumption and are able (however, unlikely) to eliminate almost all of the stochastic variation in their consumption, consumption follows a deterministic trend plus a very small unit root component. In that case $t\hat{a}y_t$ is an almost perfect proxy $c\hat{a}y_t$ and thus embodies substantial economic content.

The paper is structured as follows. In the next section, we derive the expression that relates private dissaving to expected future real asset returns and consumption and non-asset income growth. Section 3 estimates $c\hat{a}y_t$ as a cointegrating relationship between consumption, non-asset income and wealth and examines the dynamics of short-run adjustment in this system. Section 4 investigates the predictive content of $c\hat{a}y_t$ and Section 5 concludes.

2. Theory

This section derives the following exact linear relation:

$$cay_t = E_t \sum_{i=1}^{\infty} \rho_y^i \left[\omega r_{a,t+i} + (1-\omega)\Delta y_{t+i} - \Delta c_{t+i} \right] \qquad (1)$$

The notation in equation (1) is as follows: $cay_t = c_t - (1-\omega)y_t - \omega a_t$ is the log approximation to (the current value of) private dissaving that was introduced by Lettau and Ludvigson (2001), c_t is the log of private consumption C_t, a_t is the log of private net wealth A_t, y_t is the log of (after-tax) non-asset income Y_t, $r_{a,t}$ is the log of the gross real return on private assets $1 + R_{a,t}$, ω is a constant with the interpretation of the average ratio of private net wealth to the present value of private consumption, ρ_y is a number slightly less than one, E_t is the expectations operator conditioned on current information, and Δ is the backward difference operator.

The interpretation of equation (1) is that rises in private dissaving anticipate either rises in future non-asset income or the real return on private assets, or falls in future consumption. In other words, the representative agent dissaves if future non-asset income is expected to grow, or the real return on private assets is expected to increase or future consumption is expected to fall. If future consumption and non-asset income are not expected to vary much relative to the future real return on private assets, private dissavings will act to signal the representative agents' expectations of future asset returns. In that case, private dissavings should be a good predictor of future asset returns.

After-tax non-asset income appears explicitly in equation (1) because we do not define aggregate wealth to include the market value of tradeable human capital, unlike Lettau and Ludvigson (2001). Nevertheless, equation (1) is formally equivalent to their equation (9) (p. 820) if the change in non-asset income is replaced by the (unobservable) return on human capital. An advantage of equation (1) is that it can be related to earlier work on the permanent income hypothesis. Under the assumption that the real return on

private assets is constant, the permanent income hypothesis implies that the representative agent completely smoothes consumption so that consumption is not expected to change. In that case, equation (1) reduces to the log-linear version of the exact linear relation derived by Campbell (1987) which shows that increases in private dissaving signal only potential changes in after-tax non-asset income[2].

We now proceed to derive equation (1) using the log-linearization approach of Campbell and Mankiw (1989). In the derivation, all unimportant linearization constants are omitted. A transversality condition $\lim_{i\to\infty} \prod_{k=1}^{i}[1/(1+R_{a,t+k}]A_{t+i} = 0$ is sufficient to ensure that the present value of private consumption Γ_t is equal to the present value of after-tax non-asset income Φ_t plus private net wealth:

$$\Gamma_t = \Phi_t + A_t \tag{2}$$

where

$$\Gamma_t \equiv C_t + \sum_{i=1}^{\infty} \prod_{k=1}^{i}[1/(1+R_{a,t+k}]C_{t+i} \tag{2a}$$

and

$$\Phi_t \equiv Y_t + \sum_{i=1}^{\infty} \prod_{k=1}^{i}[1/(1+R_{a,t+k}]Y_{t+i}\,. \tag{2b}$$

Equation (2a) implies the following law of motion:

$$\Gamma_{t+1} = (1+R_{a,t+1})(\Gamma_t - C_t) \tag{3}$$

Divide both sides of equation (3) by Γ_t and manipulate the result to get the log-linearized difference equation:

$$\gamma_{t+1} - \gamma_t = r_{a,t+1} + \ln[1-\exp(c_t - \gamma_t)] \tag{4}$$

where $\gamma_t = \ln \Gamma_t$ and $r_{a,t+1} = \ln(1 + R_{a,t+1})$. Linearize the right-hand side of equation (4) to get:

$$\gamma_{t+1} - \gamma_t = r_{a,t+1} + (1 - 1/\rho_c)(c_t - \gamma_t) \tag{5}$$

where $\rho_c \equiv 1 - \exp(c - \gamma) = 1 - C/\Gamma$ is a number slightly less than one. The unsubscripted terms C and Γ can be interpreted as average values of C_t and Γ_t. Equation (5) can be written as the first order difference equation:

$$(c_{t+1} - \gamma_{t+1}) - 1/\rho_c (c_t - \gamma_t) = -r_{a,t+1} + \Delta c_{t+1} \tag{6}$$

Solve in the forward direction (since $1/\rho_c > 1$) to get:

$$c_t - \gamma_t = \sum_{i=1}^{\infty} \rho_c^i (r_{a,t+i} - \Delta c_{t+i}) \tag{7}$$

The next step is to derive the analogous equation for non-asset income. By symmetry with equations (3) to (7), it follows that:

$$y_t - \phi_t = \sum_{i=1}^{\infty} \rho_y^i (r_{a,t+i} - \Delta y_{t+i}) \tag{8}$$

where $\phi_t \equiv \ln \Phi_t$ and $\rho_y \equiv 1 - \exp(y - \phi) = 1 - Y/\Phi$ is another number slightly less than one. Finally, assume $A_t > 0$ for $t = 0,1,2,\ldots$, write equation (2) as

$$\phi_t - \gamma_t = \ln[1 - \exp(a_t - \gamma_t)],$$

and linearize the right-hand side of this expression to get

$$\phi_t - \gamma_t = -\frac{\omega}{1 - \omega}(a_t - \omega_t) \tag{9}$$

where $\omega = A/\Gamma$. In equation (9), use (7) to eliminate γ_t and (8) to eliminate ϕ_t and assume $\rho_c = \rho_y$, to arrive at[3]:

$$cay_t = \sum_{i=1}^{\infty} \rho_y^i \left[\omega r_{a,t+i} + (1-\omega)\Delta y_{t+i} - \Delta c_{t+i} \right]. \qquad (10)$$

Linearly project both sides of equation (10) onto the information set at time *t*, which includes, at least cay_t, in order to complete the derivation of equation (1).

3. Consumption, Wealth and Non-asset Income

(i) *Data*

The empirical analysis uses data on consumption, after-tax non-asset income and wealth that was originally constructed by Tan and Voss (2003). A full description of the method of construction of these series can be found in the appendix of their Reserve Bank of Australia Research Discussion Paper (2000). An updated and revised version of the Tan and Voss data set was kindly provided to us by the Research Department of the Reserve Bank of Australia. The series are household non-durable consumption, household after-tax non-asset income and household net wealth. Durable consumption is excluded from the consumption measure since it represents replacements and additions to a stock whereas in the intertemporal budget constraint consumption is a flow. Household after-tax non-asset income is defined as wages plus transfers less income and other taxes. Household net wealth is the sum of financial and non-financial wealth of households less household debt. All series are quarterly and seasonally adjusted. All are expressed in real per-capita terms by dividing by population (measured in thousands of persons) and by the implicit price deflator for household consumption of non-durables (99/00=1.00).

(ii) *Long-run Relationship*

In the empirical analysis, the sample is 1988:4 – 2002:1. Tan and Voss (2003) found evidence of a long-run relationship among the levels of real per-capita consumption, wealth and non-asset income for the sample 1988:4–1999:3 but found no

such evidence for the earlier sample 1980:1–1988:3. In view of their findings, 1988:4 was chosen as the start of the sample. In the empirical analysis, we depart from Tan and Voss and use the log-levels of the series since the exact linear relation of Section 2 involves logarithms of the variables. In other words, log real per-capita series are used in the empirical analysis.

Because household expenditure on non-durables is only part of total household consumption, we follow Lettau and Ludvigson (2001) and assume that the log of total household consumption c_t is proportional to the log of household consumption of non-durables $c_{n,t}$ so that $c_t = \lambda c_{n,t}$ where $\lambda > 1$. Thus, for our consumption measure, private dissaving cay_t is defined as $c_{n,t} - (\omega/\lambda)a_t - [(1-\omega)/\lambda]y_t$. The results of augmented Dickey-Fuller tests (not shown) reveal that each series in cay_t has a single unit root. Consequently, an estimate of cay_t can be obtained as $\hat{cay}_t = c_{n,t} - \hat{\beta}_a a_t - \hat{\beta}_y y_t$, provided this estimated relationship is a cointegrating one. To estimate this relationship, we use the single equation DOLS/DGLS procedure suggested by Stock and Watson (1993). Dynamic OLS (DOLS) involves regressing $c_{n,t}$ on a constant, a_t and y_t, and on the leads, lags and contemporaneous value of the first difference of both of these variables. With dynamic generalized least squares (DGLS), this regression is estimated with a correction for serial correlation in the residuals. In either case, the coefficient estimates on a_t and y_t can be used to construct an estimated series for cay_t. Although both methods give estimates that are asymptotically equivalent, in small samples, it seems preferable to use DGLS which models residual serial correlation explicitly.

Table 1 reports the coefficient estimates obtained from the DOLS regression with four leads and lags, together with their *t*-statistics, which are based on Newey-West standard errors to correct for any residual serial correlation. To test that the DOLS estimates yield a cointegrating relationship, the Kwiatkowski, Phillips, Schmidt and Shin (KPSS, 1992) test, which has the null of cointegration, is applied to the residuals from the DOLS regression. Critical values for the KPSS test when it is used as a residual based test for cointegration are provided by Shin (1994). On the basis of the five-percent Shin critical value, the KPSS test cannot reject the null of stationarity, implying that $c\hat{a}y_t$ constructed from the DOLS estimates is a stationary process. However, the null of non-stationary residuals cannot be rejected by the ADF test on the basis of the Phillips-Ouliaris (1990) asymptotic five-percent critical value. Nevertheless, in view of the low power of the ADF test against stationary alternatives, we argue the evidence favours cointegration on the basis of the KPSS test. It is interesting to note that the DOLS estimates are practically identical to those reported by Lettau and Ludvigson (2001) for U.S. data. Since we found evidence for cointegration and our sample size is small, we re-estimated the DOLS specification using DGLS assuming an AR(1) process for the residuals. Table 1 shows that the DGLS estimates are not much different from the DOLS estimates. The DGLS estimates $\hat{\beta}_a$ and $\hat{\beta}_y$ sum to 0.822 implying that λ has a value of 1.217. The ratio of asset holdings to the present value of consumption ω $(= \hat{\beta}_a \lambda)$ is then estimated to be 0.377. This ratio, which is analogous to the share of asset holdings in aggregate wealth, is slightly larger than the one-third estimated by Lettau and Ludvigson (2001).

(iii) *Dynamic Interactions*

Table 2 reports the results from estimating a vector-error correction (VEC) model of consumption, non-asset income and wealth. A vector-error correction representation exists by the Granger representation theorem since, as we have found, there is evidence for a long-run relationship among these variables. In particular, the adjustment of the variables in response to a deviation from the long-run relationship can be inferred from the estimated VEC model. The VEC model was estimated with two lags and an error-correction term constructed from the DGLS estimates reported in table 1. This lag length was chosen as it minimized the AIC criterion and was also the optimal lag length based on the Sims likelihood ratio test statistic.

Table 2 reveals three interesting properties concerning the interactions between consumption, non-asset income and wealth. First, the coefficient on the error-correction term $c\hat{a}y_{t-1}$ is small and not statistically significant in the equation for the change in consumption or non-asset income. Both consumption and non-asset income are weakly exogenous and neither variable adjusts to deviations in the long-run cointegrating relationship. In terms of the exact linear relation of Section 1, this suggests that $c\hat{a}y_t$, which measures private dissavings, is of limited value in signaling future changes in consumption or non-asset income. Since $c\hat{a}y_t$ does not predict changes in non-asset income, the data do not appear to be consistent with the version of the permanent income hypothesis that assumes a constant rate of return on private assets. Second, while consumption is weakly exogenous, it is not strictly a random walk process since lagged changes in wealth appear to predict next periods change in consumption. Third, the coefficient on the error-correction term is large and statistically significant in the equation

for the change in wealth. It is wealth and not consumption or non-asset income that is forecast to adjust to bring about a return to the long-run cointegrating relationship following a departure from it due to shocks in consumption, non-asset income or wealth. In terms of the exact linear relation, the VEC results suggest that $c\hat{a}y_t$ is potentially a predictor of changes in asset returns since it predicts changes in wealth but not changes in consumption or non-asset income. We will argue in the next section that $c\hat{a}y_t$ predicts changes in wealth precisely because it predicts changes in asset returns, consistent with the theory underlying the exact linear relation of Section 2.

It is of some interest to investigate the implied response of consumption to a wealth shock. Under an assumed recursive ordering of the variables (and shocks), a full set of impulse response functions can be obtained from the estimated VEC model once it is re-parameterized to the levels of the variables. Part (a) of Figure 1 shows the response of consumption to a one-standard deviation wealth shock when consumption is ordered first, followed by non-asset income and then wealth. For this ordering, consumption does not respond contemporaneously (that is, within the current quarter) to shocks to non-asset income or wealth. In response to a wealth shock, consumption rises by around two and a half percent after five quarters and then falls so that in the long-run consumption is unchanged. For this ordering, a wealth shock raises the level of consumption temporarily in the short-run but has no long-run effect on the level of consumption. The explanation for this result lies in the response of wealth to a wealth shock. The effect of a wealth shock on the level of wealth is transitory: wealth rises in the short-run (by around one and a half percent after three-quarters) but remains unchanged over the long term.

Part (b) of the figure shows the response of consumption to a one-standard deviation wealth shock when wealth is placed first in the recursive order followed by non-asset income and then consumption. For this ordering, consumption responds contemporaneously to all three shocks. In response to a wealth shock, consumption rises by over five percent after eight quarters and subsequently converges to its long-run response of about four percent. Here a wealth shock permanently raises the level of consumption. Again the response of wealth to a wealth shock provides an explanation. Here the effect of a wealth shock is permanent: the level of wealth rises by over two percent after three-quarters and by just over one percent in the long term.

Alternatively for VEC systems, it is possible to decompose the shocks into those that have permanent effects on the levels of the series and those that have only transitory effects. In our system, where there is only one cointegrating vector among the three series, there are two permanent shocks and one transitory shock. The moving average representation of the VEC model is:

$$\Delta X_t = C(L)\varepsilon_t \tag{11}$$

where $C(L)$ is a lag polynomial of infinite order. There are several ways to decompose the reduced-form shocks into those that have permanent and transitory effects. Gonzalo and Ng (2001) provide one such way. Define:

$$G = \begin{bmatrix} \alpha_{\perp}^{\prime} \\ \beta^{\prime} \end{bmatrix} \tag{12}$$

where $\alpha_{\perp}$ is the orthogonal complement of α, the vector of coefficients on the error-correction term in our three-variable VEC system and β is the vector of cointegrating coefficients. They show that transforming the reduced-form shocks as $G\varepsilon_t$ produces a

permanent-transitory decomposition in the sense that the last shock does not have a permanent impact on the levels of consumption, non-asset income and wealth.[4] However, the resulting shocks are not necessarily orthogonal. To obtain permanent and transitory shocks that are orthogonal to one another, transform the reduced-form shocks as $H^{-1}G\varepsilon_t$ where H is the lower triangular matrix such that $HH' = \text{cov}(G\varepsilon_t)$.

Figure 2 shows the responses of the levels of consumption, non-asset income and wealth to a one-standard deviation transitory shock that is orthogonal to the two permanent shocks. An increase in the transitory shock leads to a pronounced increase in wealth and a much smaller increase in consumption and non-asset income at short horizons. The effect of the transitory shock on wealth is particularly strong for the first year and then dissipates. This effect is less persistent than found in U.S. data by Lettau and Ludvigson (2003) where the effect of the transitory shock on wealth is strongly significant for periods up to several years. Finally, the figure shows that the strong response of wealth to the transitory shock primarily causes $c\hat{a}y_t$ to deviate from its mean which then forecasts a subsequent adjustment in wealth to restore the cointegrating relation.

4. The predictive content of $c\hat{a}y$

(i) *Predictability of Stock Returns*

In section 3, it was found that $c\hat{a}y_t$ has predictive power for changes in next period's wealth but not for consumption or non-asset income. A possible explanation for the ability of $c\hat{a}y_t$ to predict changes in wealth is that it also predicts changes in asset returns, as the exact linear relation of section 2 implies. To investigate this explanation, table 3 reports the results of predictive variable regressions for stock returns as well as

excess returns. Stock returns are quarterly real returns on the ASX200 accumulation price index. Excess returns are the quarterly real returns in excess of the quarterly real interest rate on 90-day bank accepted bills. The stock return and interest rate data were obtained from the *Statistical Bulletin* of the Reserve Bank of Australia. To calculate real returns and the real interest rate, the implicit price deflator for non-durables consumption was used. For all of the regressions in table 3, a Newey-West correction (Newey and West, 1987) is made to the *t*-statistics for generalized serial correction of the residuals.

Panel A of the table shows the results for real returns. There is no explanatory power in the regression that includes just lagged returns. However, the explanatory power of the regression increases to seven percent when lagged $c\hat{a}y_t$ is included as well or when lagged $c\hat{a}y_t$ appears by itself apart from the constant. In both regressions where lagged $c\hat{a}y_t$ appears, its coefficient is positive and statistically different from zero. Thus, when private dissaving falls, real stock returns over the following quarter are predicted to fall and vice-versa. These results accord well with the exact linear relation. If returns are expected to fall in the future, investors who desire to smooth consumption over their lifetimes, will temporarily reduce consumption below its long-term relationship with both non-asset income and wealth in an attempt to maintain future consumption in the light of lower future returns and vice-versa. The temporary reduction in consumption in response to an expected fall in returns reduces $c\hat{a}y_t$. Thus, the theory predicts a positive relationship between $c\hat{a}y_t$ and future returns, which is exactly what we find from the regression evidence.

The results for excess real returns, shown in panel B of the table, are essentially no different to those for real returns. In each case, the estimated coefficient on lagged

$c\hat{a}y_t$ is positive and statistically significant while the estimated coefficient on lagged excess returns is not statistically significant. Moreover, the explanatory power of a regression that includes lagged $c\hat{a}y_t$ as a regressor is considerably higher in the case of excess returns. Comparing equation (3) with (6) in the table, for example, the adjusted *R*-squared increases from seven to twelve percent. These results show that lagged $c\hat{a}y_t$ has predictive power for the risk premium on stocks, or in other words, for time variation in the reward for holding risky stocks over risk-free bills.

(ii) *Long-horizon Predictability*

The exact linear relation of Section 2 has the implication that $c\hat{a}y_t$ should have long-horizon predictive power for either asset returns, consumption growth or non-asset income growth. In this section, we examine the information content of $c\hat{a}y_t$ for asset returns over long holding periods and for consumption and non-asset income growth over long horizons. To investigate long-horizon predictability in a single equation framework, we would regress, for example, the *H*-period log real return, defined as $sr_{t+1} + sr_{t+2} + \ldots + sr_{t+H}$, on $c\hat{a}y_t$ for various horizons *H*. Fama and French (1988) estimate regressions of this form. However, as Hodrick (1992) and others have pointed out, in regressions of this form, the coefficient estimates and tests statistics are biased in small samples and this bias is more pronounced when the forecast horizon is large relative to the sample size. This consideration is particularly relevant here since our sample consists of only thirteen years of quarterly data.

An alternative to the single equation approach is to impute long-horizon statistics from a VAR rather than estimating them directly. This approach avoids the small sample biases that may arise from single equation regressions of the Fama and French type. The

methodology for measuring long-horizon statistics by estimating a VAR was developed by Kandel and Stambaugh (1989), Campbell (1991) and Hodrick (1992). Because our sample is small, we will follow their approach here. Details of the approach can be found in those articles and only a heuristic description will be provided here.

Consider a first-order VAR: $Z_{t+1} = AZ_t + u_{t+1}$ where A is the coefficient matrix and u_{t+1} is the vector of residuals[5]. Because the series in Z_t are assumed to be covariance stationary, the unconditional variance of the Z_t process is finite and can be calculated as an approximation to an infinite sum[6]. Next, note that the forecast of Z_{t+H} from the first-order VAR, conditional on time *t* information, is $\hat{A}Z_t$ where $\hat{A}$ is the estimated coefficient matrix. Using this result and an estimate of the unconditional variance, the unconditional autocovariances of Z_t can be calculated. All long horizon statistics, for example, the *R*-squared statistics, are functions of the unconditional variance and autocovariances of the Z_t process. These, in turn, are functions of the estimated VAR coefficients and the estimated variance-covariance matrix of the residuals. The formulae are given in Hodrick (1992).

Part A of Table 4 shows the long horizon *R*-squared statistics calculated from three first-order VAR models[7]. For the first model, $Z_t = (\Delta c_t \ \Delta y_t \ c\hat{a}y_t \ sr_t)$, for the second, Z_t is the same except that excess real returns (esr_t) replace real returns (sr_t) and, for the third, $c\hat{a}y_t$ is excluded altogether. In each case, Z_t is covariance stationary since $c\hat{a}y_t$ characterizes a cointegrating relationship and the other series are found to be stationary on the basis of ADF unit root tests (not shown). The long horizon *R*-squared statistics are often referred to as "implied *R*-squared" statistics since they are calculated

directly from the VAR. The implied R^2 statistics for horizons up to thirty-two quarters are reported for all the series in the first model and for only excess returns in the other two. This is because the results for the other series were practically identical across models.

Part A of the table shows that the VARs have no forecasting power for either future consumption growth or growth in non-asset income at any horizon since the implied R^2 statistics are all very close to zero. By contrast, the VARs that include $c\hat{a}y_t$ have significant forecasting power for future stock returns. The forecasting power is somewhat larger for the VAR with excess returns than real returns and, in both cases, the forecasting power is strongest at short to intermediate horizons. The predictive power of the VAR for returns and excess returns, respectively, is hump-shaped and peaks around one-year. For excess returns, the implied *R*-squared is as much as nineteen percent for a holding period of four quarters. For very long holding periods, the implied *R*-squared is low; for a holding period of eight years, the statistic is only six percent for excess returns. These results indicate that VARs that include $c\hat{a}y_t$ predict returns over short and intermediate horizons. The explanatory power of the VARs for returns can be attributable to $c\hat{a}y_t$, since when that variable is excluded as in the third model, the implied *R*-squared statistics for excess returns are low at all horizons, never exceeding two percent. We can conclude that it is not lagged returns but rather $c\hat{a}y_t$ that has predictive content for returns. Our results agree remarkably well with the findings of Lettau and Ludvigson (2001) for the U.S. who report that the predictive power of $c\hat{a}y_t$ for excess returns is hump-shaped and peaks at five quarters with an $\overline{R}^2$ of twenty-one percent (p. 839).

Many other variables have been found to forecast returns in the U.S. These include the log dividend-price ratio (Fama and French (1988)) and the log dividend-earnings ratio (Lamont (1998)), where the latter is also known as the payout ratio. Part B of Table 4 reports the implied *R*-squared statistics for excess returns from two VAR models, one that includes these variables along with $\hat{cay}_t$ and one that does not. Data on the dividend yield and the price-earnings ratio on the All Ordinaries index, which subsequently became the ASX200, were obtained from the Reserve Bank of Australia *Statistical Bulletin*. The payout ratio is constructed by multiplying these series together. For the model that does include $\hat{cay}_t$, the implied *R*-squared statistics are somewhat larger and have the same hump-shaped pattern as those reported in part A for the second model. This suggests that $\hat{cay}_t$ is the source of the predictive power of the VAR for excess returns. This is confirmed by the low *R*-squared statistics from the VAR that includes only the dividend yield and the payout ratio.

(iii) *Out-of-Sample Forecast Comparisons*

Since the coefficients in $\hat{cay}_t$ are estimated from the full sample, $\hat{cay}_t$ contains full sample information and this fact may partly account for its success in predicting future returns. As Lettau and Ludvigson (2002) convincingly argue, this is not a concern when testing the theoretical framework of forward looking consumers, since such a test requires the most accurate estimates of the cointegrating coefficients available and these are only obtained from the full sample. However, it is of concern for practitioners who would like to know whether $\hat{cay}_t$ has predictive content for stock returns in real time. In view of this concern, it is of some interest to forecast returns by re-estimating the coefficients in $\hat{cay}_t$ every period, using only data available at the time of the forecast.

The difficulty with this approach is that consistent estimation of the coefficients in $c\hat{a}y_t$ requires a large number of observations. Even for the full sample, the number of observations is small here. Thus, it is particularly likely that there will be significant sampling error in the estimated coefficients of $c\hat{a}y_t$ during the early re-estimations. This would make it more difficult for $c\hat{a}y_t$ to display forecasting power and thus more difficult to validate the exact linear relation when it is consistent with the data. In view of this reservation, we will consider two forecasting exercises where, in the first one, the coefficients in $c\hat{a}y_t$ are re-estimated each period while, in the second, are set to their full sample estimates.

We compare one-quarter-ahead forecasts of returns from a first-order VAR comprising returns and $c\hat{a}y_t$ with the forecasts from, respectively, an autoregressive model and a constant model of returns. Here returns are taken to mean real returns or excess real returns on the ASX200 accumulation index. In this framework, the VAR nests both alternatives since it alone includes $c\hat{a}y_t$ and thus both alternatives can be viewed as restricted models relative to the VAR. The results of the forecast comparisons are shown in table 5. In panel A of the table, we report forecast comparisons when $c\hat{a}y_t$ is estimated with data that is only available at the time the forecast is made. We first estimate $c\hat{a}y_t$ with data from 1988:4 to 1999:1. This is the smallest sample for which the estimated coefficients in $c\hat{a}y_t$ appear to resemble their full sample counterparts. In particular, for samples that end earlier, the estimated coefficient on wealth is often implausibly small and appears to be imprecisely estimated. We then estimate the VAR comprising $c\hat{a}y_t$ and returns and the two restricted models over this sample. Each model is then used to

generate a forecast of returns for 1999:2. The sample is then extended by one quarter and $c\hat{a}y_t$, as well as the VAR and restricted models are re-estimated over this sample. The models are used to forecast returns for 1999:3. By proceeding in this way, one-quarter-ahead forecasts are obtained for 1999:2-2002:1 inclusive. The forecast performance of $c\hat{a}y_t$ is measured by the ratio of the mean squared forecast error in predicting next quarter's returns from the VAR relative to that from one of the restricted models. A number less than one indicates that inclusion of $c\hat{a}y_t$ improves forecast performance. Panel B of the table reports forecast comparisons made in exactly the same way, except that $c\hat{a}y_t$ is estimated only once using the full sample.

When $c\hat{a}y_t$ is re-estimated, there is a gain of about eleven percent in forecast accuracy from the VAR relative to the AR alternative in predicting excess returns and around six percent in predicting returns. However, against the alternative of constant returns, there appears to be no gain in forecast accuracy from the VAR in the case of excess returns and a slightly worse performance in the case of returns. On the other hand, when $c\hat{a}y_t$ is estimated from the full sample, there is a significant gain in forecast accuracy from the VAR relative to the AR alternative in predicting both returns and excess returns and, to a lesser extent, against the constant returns alternative. Panel B of the table shows that the gain relative to the AR alternative is between twelve and eighteen percent and is between seven and twelve percent relative to the constant returns alternative, in terms of relative mean-squared error. In summary, the results provide strong evidence for the model when $c\hat{a}y_t$ is estimated over the full sample but are

somewhat inconclusive when $c\hat{a}y_t$ is re-estimated period by period, perhaps because the sample sizes are quite small in the early re-estimations of $c\hat{a}y_t$.

5. Conclusion

Under the assumption that the representative agent is intertemporally solvent, we derive an expression that relates private dissaving to expected future changes in consumption, non-asset income and the real return on private assets. Empirically, we find that private dissaving has predictive content for both real stock returns and excess stock returns over short and intermediate horizons. This result for excess returns implies that the risk premium on stocks is time varying. Private dissaving is also found to have negligible predictive content for changes in consumption and non-asset income at all forecast horizons, which is consistent with the permanent income hypothesis. The measure of private dissavings takes the form of a long-run cointegrating relationship between consumption, non-asset income and wealth. We find that in response to deviations from the long-run cointegrating relationship or, equivalently, to movements in private dissaving, it is wealth that adjusts and not consumption or non-asset income. It is this feature of the adjustment process that accounts for ability of private dissaving to signal future changes in real stock returns.

An important implication of our analysis is that expected future declines in stock returns need not result in concomitant declines in consumption. When stock returns are expected to fall, agents will respond by decreasing consumption out of current wealth and non-asset income. As a result, the current level of private dissavings falls thereby enabling agents to maintain a relatively smooth consumption path into the future even though stock returns are expected to fall.

References

Brennan, Michael, and Yihong Xia. "*tay's* as good as *cay*." Unpublished paper, Anderson School of Business, University of California, Los Angeles (2002).

Campbell, John Y., "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis." *Econometrica* 55 (1987a), 1249-73.

Campbell, John Y., "Stock Returns and the Term Structure." *Journal of Financial Economics* 18 (1987b), 373-99.

Campbell, John Y., " A Variance Decomposition for Stock Returns." *Economic Journal* 101 (1991), 157-79.

Campbell, John Y., and N. Gregory Mankiw. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." In *National Bureau of Economic Research Macroeconomics Annual 1989* edited by Oliver J. Blanchard and Stanley Fischer. MIT Press, Cambridge, MA.

Cochrane, John H., "Production-Based Asset Pricing and the Link between Stock Returns and Economic Fluctuations" *Journal of Finance* 46 (1991), 207-34.

Fama, Eugene, and Kenneth French. "Dividend Yields and Expected Stock Returns." *Journal of Financial Economics* 22 (1988), 3-26.

Gonzalo, Jesus, and Serena Ng. "A Systematic Framework for Analyzing the Dynamic Effects of Permanent and Transitory Shocks." *Journal of Economic Dynamics and Control* 25 (2001) 1527-46.

Hall, Robert. "Stochastic Implications of the Life-Cycle Permanent Income Hypothesis: Theory and Evidence" *Journal of Political Economy* 96 (1978), 339-57.

Hodrick, Robert J., "Dividend Yields and Expected Stock Returns: Alternative procedures for Inference and Measurement." *Review of Financial Studies* 5 (1992), 357-86.

Kandel, Shmuel, and Robert F. Stambaugh. "Modeling Expected Stock Returns for Long and Short Horizons." Working Paper 42-88, Rodney L. White Center, Wharton School, University of Pennsylvania (1989).

Kwiatkowski, Denis., Phillips, Peter C. B., Schmidt, Peter., and Yongcheol Shin. "Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root: How Sure are We that Economic Time Series have a Unit Root." *Journal of Econometrics* 54 (1992), 159-78.

Lamont, Owen. "Earnings and Expected Returns." *Journal of Finance* 53 (1998), 1563-87.

Lettau, Martin, and Sydney Ludvigson. "Consumption, Aggregate Wealth and Expected Stock Returns." *Journal of Finance* 56 (2001), 815-849.

Lettau, Martin, and Sydney Ludvigson. "*tay's* as good as *cay*: Reply." Unpublished manuscript, Stern School of Business, New York University (2002).

Lettau, Martin, and Sydney Ludvigson. "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption." National Bureau of Economic Research, Working Paper 9848 (2003).

Newey, Whitney K., and Kenneth D. West. "A Simple Positive Semidefinite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (1987), 703-8.

Phillips Peter C. B., and Sam Ouliaris. "Asymptotic Properties of Residual Based Tests for Cointegration." *Econometrica* 58 (1990), 165-93.

Shin, Yongcheol. "A Residual – Based Test of the Null of Cointegration against the Alternative of No Cointegration." *Econometric Theory*, 10 (1994), 91-115.

Stock, James H., and Mark W. Watson. "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems." *Econometrica*, 6 (1993), 783-820.

Tan, Alvin, and Graham Voss. "Consumption and Wealth in Australia" *Economic Record*, 79 (2003), 39-56.

Tan, Alvin, and Graham Voss. "Consumption and Wealth" Reserve Bank of Australia Research Discussion Paper, Number 9 (2000).

Table 1
Estimates of the Cointegrating Relationship
Sample 1988:4-2002:1

Cointegrating Relationship	$cay_t = c_{n,t} - \beta_a a_t - \beta_y y_t$		
DOLS Estimates (Lags=2)	$\hat{\beta}_a = 0.294$	$\hat{\beta}_y = 0.598$	
	(7.45)	(5.41)	
KPSS Statistic (Lags=2) = 0.21*			
KPSS Statistic (Lags=4) = 0.17*			
ADF Statistic (Lags=1) = –3.09			
DGLS Estimates (Lags=2)	$\hat{\beta}_a = 0.310$	$\hat{\beta}_y = 0.512$	$\rho = 0.79$
	(5.12)	(3.29)	(8.26)

Notes to table 1: This table reports both dynamic ordinary least squares (DOLS) and dynamic generalized least squares (DGLS) estimates of the parameters in the cointegrating relation $cay_t = c_{n,t} - \beta_a a_t - \beta_y y_t$, where c_{nt} is consumption of non-durables and services, a_t is household net wealth and y_t is non-asset income. All series are logged and expressed in real per-capita terms. Two leads and lags of the right-hand side variables (a_t and y_t) are used in the DOLS and DGLS regression. Numbers in parentheses below the coefficient estimates are *t*-statistics calculated from Newey-West (1987) standard errors with two autocovariance lags. Critical values for the KPSS test when it is used as a residual based test for cointegration are provided by Shin (1994). The appropriate five-percent critical value is 0.221. This corresponds to a cointegrating relationship involving two right-hand side variables and a DOLS specification that includes a constant but no trend. In calculating the KPSS test statistic, we use the weighting scheme of Newey and West (1987) with the number of autocovariance lags as indicated. The ADF statistic is the augmented Dickey-Fuller *t*-statistic for the null of a unit root in the residuals from the DOLS regression. The ADF regression includes one lagged change in the residuals as that lag length was selected by the AIC criterion. Phillips and Ouliaris (1990) provide critical values for the ADF test when it is used as a residual based test for the absence of cointegration. The appropriate asymptotic five-percent critical value is –3.77. An asterisk denotes significance at the five-percent level.

Table 2
Estimates of Vector-Error Correction Model
Sample 1988:4 – 2002:1

Dependent Variable	Equation		
	Δc_t	Δa_t	Δy_t
$\sum_{i=1}^{2} \Delta c_{t-i}$	–0.208	0.006	0.250
(*t-stat*)	(–0.82)	(0.01)	(0.54)
$\sum_{i=1}^{2} \Delta a_{t-i}$	**0.161**	0.434	0.263
(*t-stat*)	(2.12)	(1.80)	(1.90)
$\sum_{i=1}^{2} \Delta y_{t-i}$	0.197	–0.076	0.041
(*t-stat*)	(1.65)	(–0.20)	(0.19)
$c\hat{a}y_{t-1}$	–0.012	**0.439**	–0.009
(*t-stat*)	(–0.18)	(2.03)	(–0.07)
$\overline{R}^2$	0.08	0.28	0.04

Notes to table 2: This table reports the coefficient estimates for a vector error-correction model. $c\hat{a}y_t = c_{n,t} - \hat{\beta}_a a_t - \hat{\beta}_y y_t$ where $\hat{\beta}_a$ and $\hat{\beta}_y$ are the DGLS estimates reported in table 1. The estimated coefficient on $c\hat{a}y_{t-1}$ in each equation is reported, together with its *t*-statistic. In each equation, the sum of the estimated coefficients on the lags of each variable are reported, together with the *t*-statistic for the sum. Significant coefficients at the five-percent level are highlighted in bold face. Also shown is the adjusted *R*-squared statistic.

Table 3
Quarterly Forecasting Regressions
Sample 1988:4 – 2002:1

Equation **Number**	*Constant* (*t-stat*)	$DepVar_{t-1}$ (*t-stat*)	$c\hat{a}y_{t-1}$ (*t-stat*)	$\overline{R}^2$	*DW*
		Panel A: Real Returns			
1	**0.020**	–0.092		–0.01	2.02
	(2.24)	(–0.62)			
2	**0.976**		**1.197**	0.07	2.27
	(2.12)		(2.08)		
3	**1.064**	–0.146	**1.304**	0.07	2.00
	(2.21)	(–1.06)	(2.17)		
		Panel B: Excess Real Returns			
4	0.007	–0.070		–0.02	2.02
	(0.84)	(–0.46)			
5	**1.212**		**1.507**	0.11	2.33
	(2.77)		(2.75)		
6	**1.354**	–0.168	**1.684**	0.12	2.02
	(2.96)	(–1.24)	(2.93)		

Notes to table 3: The table reports the coefficient estimates from OLS regressions of real stock returns and excess real stock returns on the variables named at the head of a column. Real stock returns are defined as $\log(ASX / IPDND)_t - \log(ASX / IPDND)_{t-1}$ where *ASX* is the ASX200 accumulation stock price index and *IPDND* is the implicit price deflator for non-durable goods. Log is the natural logarithm. Excess real stock returns are obtained by subtracting the real interest rate on 90-day Bank Accepted Bills from real returns. The real interest rate is defined as $\log(1 + BB90_t / 400) - \log(IPDND_t / IPDND_{t-1})$ where *BB90* is the interest rate on 90-day bank bills, expressed as percent per year. The one-period lag of the dependent variable is denoted by $DepVar_{t-1}$. The variable $c\hat{a}y_t$ is defined as in table 2. Numbers in parentheses below the coefficient estimates are *t*-statistics calculated from Newey-West (1987) standard errors based on two autocovariance lags. Significant coefficients at the five-percent level are highlighted in bold face. Also shown is the adjusted *R*-squared and the Durbin-Watson (*DW*) statistic.

Table 4
Implied Long-horizon R-Squared Statistics from VAR(1) Models

	Forecast Horizon (in quarters)								
Equation	1	2	3	4	8	12	16	24	32
Part A									
(a) $Z_t = (\Delta c_t \ \Delta y_t \ c\hat{a}y_t \ sr_t)$									
Δc_t	0.04	0.02	0.02	0.01	0.01	0.01	0.00	0.00	0.00
Δy_t	0.03	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00
$c\hat{a}y_t$	0.63	0.59	0.52	0.46	0.30	0.20	0.15	0.09	0.07
sr_t	0.08	0.10	0.11	0.12	0.12	0.10	0.08	0.06	0.04
(b) $Z_t = (\Delta c_t \ \Delta y_t \ c\hat{a}y_t \ esr_t)$									
esr_t	0.12	0.15	0.18	0.19	0.18	0.15	0.12	0.08	0.06
(c) $Z_t = (\Delta c_t \ \Delta y_t \ esr_t)$									
esr_t	0.02	0.01	0.01	0.02	0.00	0.00	0.00	0.00	0.00
Part B									
(d) $Z_t = (d_t - p_t \ \ d_t - e_t \ \ c\hat{a}y_t \ \ esr_t)$									
esr_t	0.15	0.21	0.23	0.23	0.18	0.13	0.10	0.06	0.05
(e) $Z_t = (d_t - p_t \ \ d_t - e_t \ \ esr_t)$									
esr_t	0.03	0.04	0.05	0.05	0.06	0.05	0.05	0.04	0.03

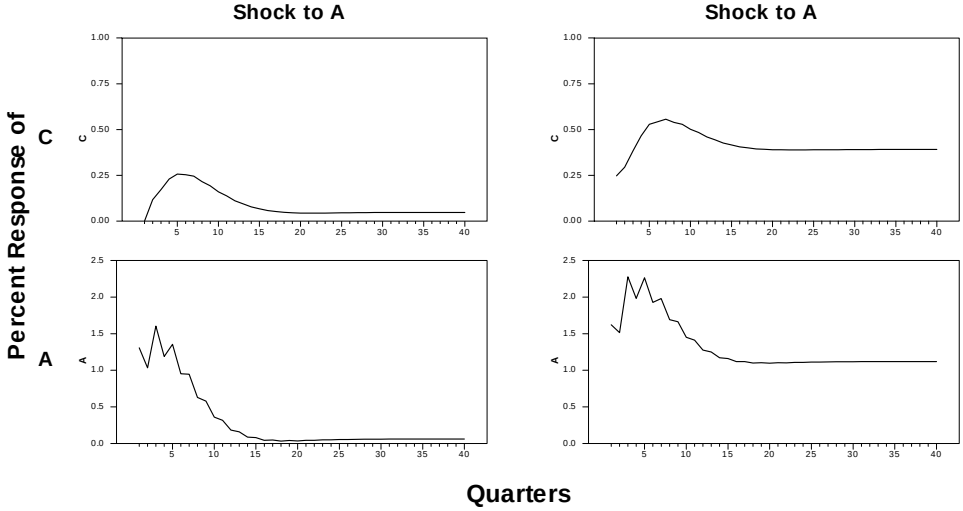
Notes to table 4: Estimated VAR models of the form $Z_t = \hat{A}Z_{t-1}$ are used to forecast Z_{t+H} for various forecast horizons *H*. The explanatory power of the VAR for a given variable in *Z* at horizon *H* is reported in the table as the "Implied R^2" from the VAR. The "Implied R^2" is calculated from the estimated parameters of the VAR, namely $\hat{A}$ and the estimated covariance matrix of the VAR residuals. In part (a), $Z_t = (\Delta c_t \ \Delta y_t \ c\hat{a}y_t \ sr_t)$ where c_t is (log) consumption of non-durables and services, y_t is (log) non-asset income, Δ is the first difference operator, $c\hat{a}y_t$ is estimated private dissaving and sr_t is the real return on the ASX200 accumulation stock price index. In part (b), Z_t is similarly defined except that the excess real return on the ASX200 (esr_t) replaces the real return. In part (c), $c\hat{a}y_t$ is excluded from the Z_t of part (b). In parts (d) and (e), Z_t contains the log dividend-price ratio ($d_t - p_t$) and the log dividend-earnings ratio ($d_t - e_t$). The variables in Z_t are demeaned by regression on a constant prior to entering the VAR. The implied *R*-squared statistics are reported for all the variables in Z_t in part (a) and only for esr_t in parts (b) to (e).

Table 5
Model Comparison of One-Quarter-Ahead Forecasts of Returns

Row	**Model Comparison**	MSE_u / MSE_r
	Panel A: Cointegrating Vector Reestimated	
1	VAR($sr, c\hat{a}y$) vrs AR(*sr*)	0.935
2	VAR($sr, c\hat{a}y$) vrs *const*(*sr*)	1.029
3	VAR($esr, c\hat{a}y$) vrs AR(*esr*)	0.891
4	VAR($esr, c\hat{a}y$) vrs *const*(*esr*)	0.992
	Panel B: Fixed Cointegrating Vector	
5	VAR($sr, c\hat{a}y$) vrs AR(*sr*)	0.879
6	VAR($sr, c\hat{a}y$) vrs *const*(*sr*)	0.923
7	VAR($esr, c\hat{a}y$) vrs AR(*esr*)	0.815
8	VAR($esr, c\hat{a}y$) vrs *const*(*esr*)	0.880

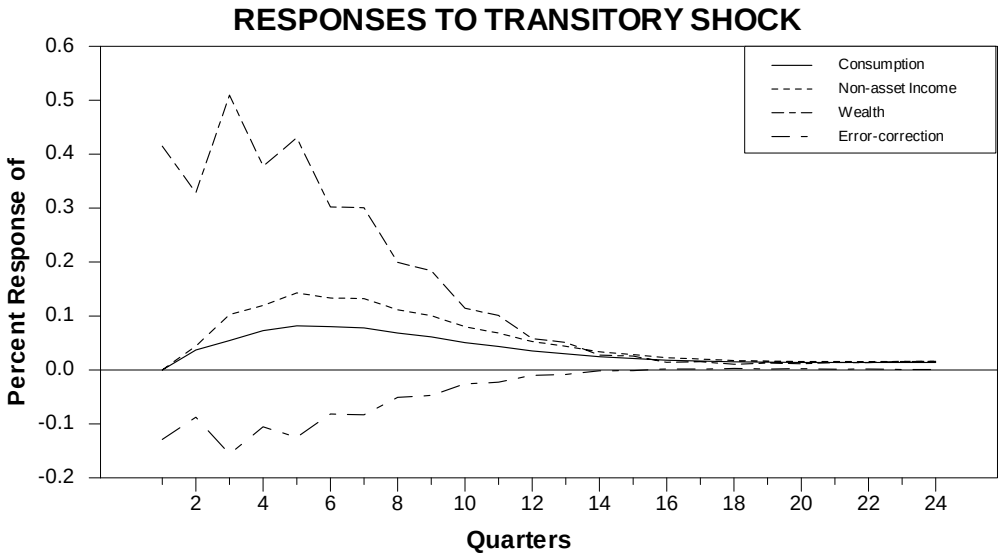
Notes to table 5: The table reports results of comparisons of one-quarter-ahead forecasts of real returns (denoted *sr*) and excess real returns (denoted *esr*) on the ASX200 accumulation stock price index. $c\hat{a}y$ denotes the estimated cointegrating vector among consumption, wealth and non-asset income. The notation VAR($sr, c\hat{a}y$) denotes a one-lag VAR model of real returns and $c\hat{a}y$ while AR(*sr*) denotes a one-lag autoregressive model for real returns. Both models include constant terms. The model where real returns are regressed on a constant only is denoted by *const*(*sr*). When the models are used to forecast excess real returns, *esr* replaces *sr* in the notation. The AR or *const* alternative is restricted relative to the VAR, which is the unrestricted model in each comparison. MSE_u / MSE_r denotes the ratio of the one-quarter-ahead mean squared forecast error from the unrestricted model (the VAR) expressed as a fraction of that from the restricted model (either the AR or *const*). In each comparison, the models are initially estimated with data from 1988:4 to 1999:1 and then used to produce one-quarter-ahead forecasts for 1999:2. The sample is then extended by one quarter, the models are re-estimated and another set of one-quarter-ahead forecasts are produced. By continuing in this way, one-quarter-ahead forecasts are generated from the models being compared for 1999:2 to 2002:1. The average of the one-quarter-ahead forecast error over this period from the VAR and the restricted model is calculated, respectively, and used to form MSE_u / MSE_r. In panel A, results are reported for the case where $c\hat{a}y$ is first estimated over the initial sample 1988:4-1999:1 and is then re-estimated each time the sample is extended by one quarter. The estimation method is DGLS with two leads and lags. In panel B, results are based on the DGLS estimate of $c\hat{a}y$ for the full sample 1988:4-2002:1, reported in table 1.

Figure 1. Response of Consumption to Wealth Shocks



(a) Order: C-Y-A (b) Order: A-Y-C

Figure 2



Notes to figure: Responses to a one-standard deviation transitory shock. The error-correction term is $\hat{cay}_t$.

Footnotes

[1] One exception is the ratio of investment to the capital stock, which Cochrane (1991) found to have some success in predicting stock returns.

[2] Specifically, under the assumption of a constant return to private assets, equation (1) (ignoring the constant term) is $cay_t = E_t \sum_{i=1}^{\infty} \rho_y^i \left[(1-\omega)\Delta y_{t+i} - \Delta c_{t+i}\right]$ and, furthermore, from the representative agent's problem, $E_t \Delta c_{t+i} = 0$ for all i (Hall, 1978). Thus, $cay_t = (1-\omega) E_t \sum_{i=1}^{\infty} \rho_y^i \Delta y_{t+i}$ which is the log-linear version of Campbell's (1987) exact linear relation.

[3] A similar assumption was also made by Campbell and Mankiw (1989). The justification for this assumption is that each discount factor is a number slightly less than one and thus approximately equal, since each factor is of the form $1 - x/PV(x)$, where $PV(x)$ is the present value of *x* over an infinite horizon.

[4] Specifically, the last column of the *3x3* long-run impact matrix $C(1)G^{-1}$ is a column vector of zeros and corresponds to the long-run impact of the last shock.

[5] Higher-order lags can be stacked to form a first-order system in which case the matrix *A* becomes (what is commonly known as) the companion matrix whose free elements are the coefficients in the VAR. See, for example, Campbell (1987).

[6] Because the series are covariance stationary, the first-order VAR can be inverted to obtain the moving average representation: $Z_{t+1} = \sum_{j=0}^{\infty} A^j u_{t+1-j}$. The unconditional variance of the Z_t process is then: $C(0) = \sum_{j=0}^{\infty} A^j V A^{j\prime}$, where $V = E(u_{t+1} u_{t+1}^{\prime})$. In the actual calculation of the long-run (i.e., unconditional) variance of Z_t, a lag truncation of 40 quarters was used.

[7] A first-order VAR was found to be adequate for all the models reported in table 4 since that lag length was selected in every case by both the AIC criterion and the Sims Likelihood ratio test.