The sensitivity of aggregate consumption to human wealth

Hasan Bakhshi

Bank of England, Threadneedle Street, London, EC2R 8AH. The views expressed are those of the author and do not necessarily reflect those of the Bank of England. I would like to thank for their help and comments: Erik Britton, Alec Chrystal, Shamik Dhar, Jens Larsen, Steve Millard, Richard Thornton, Simon Wren-Lewis, Tony Yates, and in particular Roy Cromb and Simon Price. The comments from two anonymous referees materially improved the paper.

Issued by the Bank of England, London, EC2R 8AH, to which requests for individual copies should be addressed; envelopes should be marked for the attention of Publications Group. (Telephone: 0171-601 4030). Working Papers are also available from the Bank's Internet site at http://www.bankofengland.co.uk/wplist.htm

ISSN 1368-5562

Contents

Abstract

The Permanent Income Hypothesis (PIH) assumes that individuals base their decisions on lifetime wealth, not current income. Textbook versions of the PIH predict that the elasticity of consumption with respect to human wealth is equal to the share of human wealth in total wealth. Comparing calibrated wealth shares with econometrically estimated elasticities amounts to a simple test of the PIH. In the United Kingdom, aggregate consumption is found to be more sensitive to changes in human wealth than is predicted by the PIH. This does not appear to be explained by a simple, but common, treatment of credit constraints.

1 Introduction and summary

This paper is about the theory and estimation of the consumption function. In particular, it focuses on the sensitivity of aggregate consumption to changes in human wealth. The main points are: first, that consumption is more sensitive to changes in human wealth than is predicted by standard treatments of consumption; and second, that the excess sensitivity does not appear to be explained by a simple, but common, treatment of exogenous credit constraints. The paper concludes that an endogenous treatment of credit constraints is likely to be more fruitful in generating an explanation.

Forward-looking treatments of consumption assume that individuals base their consumption decisions on the present value of expected lifetime income (human wealth), not current income. Some theories (including textbook versions of the permanent income hypothesis (PIH) and the Blanchard-Buiter overlapping generations model) imply consumption functions that generate at least two testable predictions. First, the elasticity of consumption with respect to total wealth (if this can be measured) should equal unity. Second, the elasticity of consumption with respect to human wealth should equal its share in total wealth.[1] This second prediction is tested in this paper.

If labour income follows a random walk with drift (so that *growth* in income fluctuates around a constant mean), the PIH[2] predicts that the elasticity of consumption with respect to human wealth equals the elasticity with respect to labour income. This paper shows that a random walk with drift is a reasonable description of the data. Comparing calibrated estimates of the human wealth share with econometrically estimated elasticities with respect

---

(1) I would like to thank Roy Cromb for making this point.

(2) In what follows we denote as the PIH all forward-looking models of consumption that generate a simple linear levels consumption function.

to labour income, the work reported below concludes that estimated elasticities in aggregate time series consumption functions are generally higher than the estimate of the human wealth share: consumption is more sensitive to changes in human wealth (and labour income) than is predicted by the PIH.

One commonly advanced explanation for this result is that some individuals are credit-constrained (see Campbell and Mankiw (1989, 1991) and Muellbauer (1994)). Intuitively, if such individuals are constrained to finance consumption out of their current labour income, consumption in the aggregate might be expected to be more responsive to changes in current income than in a world where all individuals can borrow and lend freely. But this explanation implicitly assumes that labour income evolves according to a specific statistical process. This paper shows that if the evolution of labour income for PIH consumers is described by a random walk with drift—and if credit-constrained consumers are assumed to have a constant proportional claim on national income—the elasticity of consumption with respect to labour income may in fact be lower, not higher, than in an economy where individuals are all unconstrained (as in the PIH). This result seems counterintuitive.

Suppose that the elasticity for individual PIH consumers is equal to the human wealth share, estimated at x, say, on aggregate data ($0 < x < 1$). The elasticity for individual credit-constrained consumers is unity by construction (they are assumed to consume all their labour income). Then if the economy is populated by a mixture of PIH and credit-constrained consumers, the paper shows that the elasticity of consumption with respect to labour income does not necessarily lie between x and unity. Such a result might not seem possible in a linear model. But the reason becomes clear when it is recognised that the elasticity for PIH consumers — the human wealth share — is endogenous. If there are credit-constrained individuals in the economy, then the implied human wealth share for PIH consumers is some number less than x. And in the case where labour income follows a random walk with drift it turns out that the weighted average of unity (the credit-constrained consumers) and this number is in general smaller than x. One must avoid the temptation of treating the elasticity for PIH consumers as a parameter, when it is in fact a variable. This is an easy mistake to make when the elasticity is estimated as a parameter in a log-linear consumption function.

The point made is not that credit constraints cannot explain the observed discrepancy between the PIH and empirically estimated elasticities of consumption with respect to labour income. It is that the particularly simplistic treatment of credit constraints that is sometimes used to account for this discrepancy will not do the trick. There are more sophisticated treatments of credit constraints in solved-out consumption functions in the literature, such as King (1986), where credit constraints are made endogenous. Such treatments are likely to be more fruitful for macro modellers who wish to estimate or identify the responsiveness of aggregate consumption to changes in labour income.

Finally, before proceeding, it is worth noting the relationship between this paper and the vast existing literature that tests for the 'excess sensitivity' of consumption. There is plenty of econometric evidence for the proposition that consumption responds to a change in labour income by more than that is warranted by an innovation in labour income in standard versions of the PIH (see the seminal piece by Flavin (1981), surveys by Deaton (1992) and Muellbauer and Lattimore (1995), and Ellwood (1998) for a recent example). Although the focus in this paper is slightly different — I look at the long-run sensitivity of consumption to innovations in human wealth — the results in this paper are consistent with the findings of that literature.

2 The model

In an infinite horizon model where individuals have a utility function with constant relative risk aversion (CRRA)[3] and they face no uncertainty about future labour income, the level of consumption of the representative agent is given by the following linear function:[4]

$$c_t^{PIH} = b\left(hw_t^{PIH} + fw_t^{PIH}\right) \qquad \textbf{(1)}$$

where $b \approx r - \frac{(r-d)}{q}\frac{1}{(1+r)(1+d)}$ and $hw_t^{PIH} = \sum_t^{\infty} \frac{yl_s^{e^{PIH}}}{(1+r)^{s-t}}$

r is the real interest rate, δ is the individual's subjective discount rate, $1/\theta$ is the intertemporal elasticity of substitution, hw_t and fw_t are human and non-human wealth respectively at time t, and yl_s^e is the individual's expected labour income at time s. The superscript *PIH* refers to the individual being a PIH consumer. As in Deaton (1992), Campbell and Mankiw (1989, 1991) and, more recently, Blake, Camba-Mendez and Weale (1998), r and d are assumed constant for simplicity. Consumption is a function of annuitised lifetime wealth (permanent income), where lifetime wealth is comprised of human wealth (discounted future labour income) and non-human wealth. The annuity rate is equal to b, where the first term, r, can be thought of as the income effect of a real interest rate change and the second term, $\frac{(r-d)}{q}\frac{1}{(1+r)(1+d)}$ is the substitution effect.[5]

---

[3] $U(c_t) = \frac{c_t^{1-q}}{1-q}$, where $1/q$ is the intertemporal elasticity of substitution.

[4] See the annex for derivation of this consumption function, and the discussion in chapter 6 of Bank of England (1999) *Economic Models at the Bank of England*.

[5] So if there is no substitution, then $b=r$, as with a conventional financial annuity with infinite life.

It is more difficult to solve for this consumption function analytically in an infinite horizon model when utility does not take the CRRA form (see Blanchard and Fischer (1989) chapter 2). And it is even more difficult when individuals face uncertainty about their future labour income. There are two special cases for which a solution is straightforward: first, when labour income risk is perfectly diversifiable and utility functions are in the hyperbolic absolute risk aversion (HARA) class;(6) second, if preferences are quadratic (as assumed in Hall (1978)).(7) Blanchard (1985) and Buiter (1988) relax the infinite horizon assumption and present continuous-time overlapping generation models (OLG) where uncertainty arises because individuals face a constant probability of death. They show that the solved-out aggregate consumption function takes the form in equation **(1)**, except that the annuity rate, b must be adjusted (upwards) to reflect the positive probability of death.

Clearly the assumptions that are required to derive the linear consumption function in equation **(1)** are neither trivial nor entirely plausible. But the tractability of the resulting aggregate consumption function and the fact that it is founded on micro-foundations means that it remains highly influential in the macro modelling literature (in both its infinite horizon or OLG forms). In fact, as noted by Deaton (1992):

'there is a sense in which the equation has a life of its own. One might simply assert, at least as a working hypothesis, that consumption is the annuity value of human wealth and non-human wealth.'

Returning to the infinite horizon model, the representative agent's consumption function in **(1)** is interpreted as representing aggregate

---

(6) The HARA class includes the CRRA, the constant absolute risk aversion (CARA) and quadratic utility functions.

(7) See Blanchard and Fischer (1989) chapter 6 for formal derivation of the consumption function in these cases.

consumption when the population of consumers has a specific age and income distribution:

$$C_t = b\left(HW_t + FW_t\right) \tag{2}$$

where C_t is aggregate consumption, and HW_t and FW_t are the aggregate stock of human and non-human wealth respectively at time *t*.

The simple levels consumption function in equation **(2)** generates at least two testable predictions. First, a regression of log consumption on log *total* wealth, controlling for the annuity rate, should yield a coefficient of unity. In simple specifications of the long-run consumption function, this restriction has been rejected on aggregate UK data (eg see Price (1999)). Second, the elasticity of consumption with respect to human wealth is equal to the share of human wealth in total wealth:

$$\frac{\partial C}{\partial HW}\frac{HW}{C} = b\,\frac{HW}{b\left(HW + FW\right)} = \frac{HW_t}{HW_t + FW_t} \tag{3}$$

(Similarly the elasticity with respect to non-human wealth is its share in total wealth).

This second prediction is tested in this paper. As we shall see, testing this particular prediction also permits us directly to consider the effect of assuming that some consumers are constrained to consume only out of their current income.

3 A test of the Permanent Income Hypothesis

3.1 Human wealth estimates

Unlike non-human wealth, human wealth is not directly observable, so estimates are hard to come by in the literature. Because of these measurement difficulties, researchers often prefer to quasi-difference[(8)] the consumption function to eliminate the term in human wealth (See Hayashi (1982) and Darby and Ireland (1994)), or to use labour income as a proxy (which would be valid if labour income follows a random walk with drift, eg in the Bank of England and HM Treasury consumption functions.

While there are some published estimates for human wealth in overseas countries, eg Macklem (1997) for Canada, there are no published estimates for the United Kingdom.[(9)]

Human wealth is equal to the discounted present value of future expected labour income:

$$HW_t = \sum_t^{\infty} \frac{YL_s^e}{(1+r)^{s-t}} \qquad \textbf{(4)}$$

I generate income forecasts, YL_s^e using a simple univariate forecasting rule for labour income.[(10)]

(8) Quasi-differencing involves subtracting discounted future consumption from current consumption.

(9) Price (1999) is an exception.

(10) I use the (real) labour income measure in the Bank's forecasting model, defined as total income from employment (including self-employment), plus (local and central) grants to the personal sector, minus employment income tax payments, poll tax payments and national insurance contributions, all deflated by the consumption deflator.

In fact, a simple random walk with drift fits the data reasonably well over our 1975 Q1 ? 1998 Q4 sample period:[11] the error term is normally distributed and serially uncorrelated, and the drift parameter, *g*, is sensible (because it conforms to trend real growth):[12]

$$dLYL = \underset{(3.49)}{0.0054} \qquad \textbf{(5)}$$

(t-statistic in brackets)

where *dLYL* denotes the growth in labour income, implying an annual growth rate of 2.2%. I investigate sensitivity of the resulting human wealth share estimates to this figure later in the robustness analysis .[13]

Assuming that individuals know the value of their labour income in the current period, equation **(5)** is used to generate forecasts for labour income at each point in time. The forecast period is truncated at twenty years in the central case. This is roughly the remaining working lifetime of the average individual, assuming a uniform age distribution and a forty-year working life. Of course, one could argue that truncating the forecast period at any time is not consistent with a strict interpretation of the infinite horizon model. But the infinite horizon assumption is never taken literally in empirical work on the PIH. Truncation can be justified by assuming that individuals place

---

(11) Price (1999) also estimates a multivariate forecasting model for labour income, but this does little to the estimated values of human wealth: the correlation between the two human wealth measures is near unity. For the United States, however, Muellbauer (1996) finds that a multivariate model for predicting US labour income is superior to a random walk with drift.

(12) Blake, Camba-Mendez and Weale (1998), using a much longer historical data set, argue that labour income is stationary and use a forecasting rule where consumers make an assumption about the rate at which income growth deviates from its trend rate.

(13) Strictly speaking I should use labour income per capita in the human wealth calculations so as to abstract from that component of the increase in human wealth over time that is due to population growth. But by the same argument I would use non-human wealth per capita to calculate the human wealth share. The population term would drop out in the share so I continue with an aggregate measure.

effectively zero weight on income streams beyond the average lifetime of a representative agent (equivalently, in the OLG model, that individuals do not have strong bequest motives). A strict interpretation of the infinite horizon assumption would also mean having to impose the restrictive stability condition that the real interest rate used to discount expected future labour income exceeds the trend growth of labour income (this is violated, for example, in the baseline case).

The rate at which individuals discount their future labour income is of course very difficult to measure. On one view the *ex ante* real rate of return on labour income is best proxied by the *ex ante* real rate of return on a risky portfolio. But this itself has huge measurement problems. An alternative is simply to use an *ex post* measure of the interest rate. This is the route followed by Darby and Ireland (1994). They used a post-tax measure of the rate of return on building society share accounts, deflated by the current rate of inflation.(14) I experimented with these different options. In view of the great measurement difficulties, I use the average post-tax real rate of return on building society share accounts between January 1980 and December 1998 (in the central case) which comes out at 2.1% per annum.(15) So the resulting time series for human wealth should probably be interpreted as a trend measure of the true series, assuming that the real interest rate is stationary.(16) Because my test of the PIH involves comparing the *average* human wealth share over the sample period with the econometrically estimated elasticity, using a constant average real interest rate should not bias the results.

---

(14) Actually, Darby and Ireland (1994) add a mark-up to the real interest rate that depends on the extent of financial deregulation.
(15) The 1970s were excluded from this calculation as *ex post* measures of the real interest rate invariably take implausibly high negative values in the 1970s (thus showing the limitations of using an *ex post* measure as a proxy for the *ex ante* real interest rate).
(16) Of course the real interest rate is a bounded variable and so is stationary by definition.

Given the potential sensitivity of the human wealth estimates to assumptions about (a) the discount rate, (b) the size of the drift term in the process for labour income, *g*, (c) the appropriate forecasting horizon for individuals, and (d) the uncertainty surrounding the appropriate values for these variables, Table A presents human wealth shares on different assumptions as a robustness check.

Table A: Estimated average human wealth shares, %, 1975 Q1 – 1998 Q4[(a)(b)]

Forecast horizon	20 years g=0.54% *pq* (=2.2% *pa*)	20 years g=0.59% *pq*[(c)] (=2.4% *pa*)	20 years g=0.74% *pq* (=3% *pa*)	30 years g=0.54% *pq* (=2.2% *pa*)
r=0.74% *pq* (=3% *pa*)	76.8	-	-	-
r=0.51% *pq* (=2.1% *pa*)	**78.4**	78.7	79.7	80.3
r=0.25% *pq* (=1% *pa*)	80.1	-	-	-

pq denotes quarterly interest rate; *pa* is the annual interest rate.

[(a)] Baseline estimate in bold. Note that $r \approx g$ in the baseline case.

[(b)] Human wealth share equals human wealth divided by human wealth plus real net financial wealth plus gross housing wealth.

[(c)]This corresponds to the assumption made in Price (1999).

Surprisingly, and reassuringly, the sensitivity of the human wealth share to assumptions about the drift term in labour income, the discount rate and the forecast horizon is not great.[(17)] So in the baseline case, the PIH would predict that the elasticity of consumption with respect to human wealth is around 0.78. Comparing this with an econometrically estimated elasticity

[(17)] Note that it is still possible to calculate positive finite human wealth shares using the formula in equation **(1)** even when the trend rate of growth of labour income exceeds the real interest rate (as in the baseline case) because the forecast horizon is truncated.

amounts to a test of the PIH. Specifically, I test the restriction that the econometrically estimated elasticity is equal to 0.78.

Arguably the fact that I use *ex post* real interest rates on building society accounts significantly understates the riskiness of future labour income and therefore the discount rate used by agents. This could significantly bias my human wealth share estimates in an upwards direction. Blake, Camba-Mendez and Weale (1998) assume a much higher real interest rate used to discount future labour income of 7.6% per year; Hayashi (1982) uses even higher estimates, ranging from 13.2% to 17.3%. So it is probably best to interpret my 0.78 estimate as an upper bound.

3.2 The elasticity of consumption with respect to human wealth

In many consumption functions (Bank of England, HMT and NIESR), the difficulties of measuring human wealth lead modellers to proxy it with labour income. This is valid in a log-linear macro model if log labour income follows a random walk with drift process. Such consumption functions yield elasticities of consumption with respect to labour income, not human wealth. But we show that for this labour income process, the two elasticities are in fact identical. This allows us to compare the human wealth share with directly estimated elasticities of consumption with respect to labour income:

$$\log YL_s = \log YL_{s-1} + g + e_s$$

This equation states that growth in labour income fluctuates around a constant trend, g.

By repeated substitution of log YL_s ,

$$\log YL_s = \log YL_0 + sg + e_s + ... + e_0$$

Taking expectations at time 0 and assuming for simplicity that $(\log YL_s)^e \approx \log\left(YL_s^e\right)$ gives:

$$\log YL_s^e \approx \log YL_0 + s\log(1+g) \Rightarrow YL_s^e = (1+g)^s YL_0 \quad \textbf{(6)}$$

using $\log(1+g) \approx g$ for small g.

Recalling the expression for human wealth:

$$HW_0 = \sum_0^\infty \frac{YL_s^{e^{PIH}}}{(1+r)^s} = \left(\sum_0^\infty \left(\frac{1+g}{1+r}\right)^s\right) YL_0 = \left(\frac{1+r}{r-g}\right) YL_0 \quad \textbf{(7)}$$

So human wealth at time 0 is a linear function of labour income at time 0 (with no additional terms). If the economy at any time t is characterised as populated by infinitely lived agents born in that period, this linear relationship will apply at all times.

As human wealth is a linear function of labour income (with no additional terms), and consumption is linear in human wealth under the PIH, the PIH predicts that the elasticity of consumption with respect to human wealth is equal to the elasticity with respect to labour income.

If cointegration techniques are applied to estimate the long-run determinants of consumption, equations **(2)** and **(7)** form the basis of the cointegrating vector as, for example, in the Bank's and HM Treasury's macro model (Bank of England (1999) and Chan, Savage and Whittaker (1995)).(18)

Estimates of the elasticity of consumption with respect to labour income in the long run, taken from the long-run consumption functions in the main macroeconometric models, are higher than the 0.78 predicted by the calibrated human wealth share:

Table B(a)

Bank of England	NIESR	Sheldon & Young	HMT
0.87	0.91	0.89	0.89

(a)These consumption functions have been estimated over sample periods not identical to the 1975 Q1-1998 Q4 sample period used in the human wealth calibration. For example the Bank's consumption function is estimated over the period 1975 Q1-1992 Q1. I re-estimated the Bank's consumption function over the sample period 1975 Q1-1998 Q4 and the estimated elasticity with respect to labour income was even greater, at 0.94.

In the Bank's consumption function we tested formally the restriction that the elasticity of consumption with respect to labour income equalled the estimates for the human wealth share shown in Table A. In all cases the restriction was easily rejected. Given the earlier argument that my human wealth share estimates may well be subject to a large upwards bias, these results suggest that consumption is significantly more sensitive to human wealth than is consistent with the PIH as set out above.(19)

---

(18) In fact, because the human wealth proxy and non-human wealth have to be estimated as separate independent variables, the consumption functions estimated are actually log approximations of the levels consumption function in **(2)**. This may impart a bias to our tests ? though it is not clear which way the bias will go (technical details of the log approximation are available from the author on request).

(19) Real interest rates are assumed to be constant in this analysis. Time-varying real interest rates may in principle bias upwards the estimated elasticity of consumption with respect to labour income. Unless a real interest rate term is included in the consumption function, movements in the real rate will be captured in the equation residual. But real interest rates also enter the right-hand side of the equation through the non-human wealth term: non-human wealth is partially revalued when real interest rates change. The correlation between non-human wealth and the residual will be negative in this case and so the OLS estimate of the elasticity with respect to non-human wealth will be biased downwards. By homogeneity, the elasticity with respect to labour income will be biased upwards. Preliminary work at the Bank of England using non-human wealth stripped of revaluation effects suggests that the bias may not, however, be great in practice.

4 Can credit constraints explain the apparent extra sensitivity?

One explanation for this result that has received support in the literature is that some consumers are credit-constrained. Intuitively, if such individuals are constrained to consume only out of their current labour income, consumption in the aggregate might be expected to be more responsive to changes in current income than in a world where all individuals are unconstrained.[20]

I show here that if labour income follows a random walk with drift, the elasticity will not in general be higher if there are credit-constrained agents in the economy.

Credit-constrained consumers are unable to borrow off their future labour income and are further assumed to hold zero net financial assets (and their housing assets are assumed to offset their mortgage liabilities).[21] So consumption by a credit-constrained consumer is given by:[22]

$$c_t^{C-C} = yl_t^{C-C} \tag{8}$$

---

[20] One argument against this is that credit constraints would not alter the long-run relationship between consumption and labour income, as over long periods, actual and permanent income move together. On this argument there is an *a priori* objection to credit constraints as an explanation for the long-run discrepancy between the PIH and the data. But against this, an economy might be viewed as having imperfect capital markets even in steady state.

[21] Darby and Ireland (1994) assume that credit constrained consumers hold net negative assets. Another alternative would be to assume that credit-constrained individuals hold positive net assets as a buffer against future shocks to income.

[22] This is an ad hoc, yet common, assumption in the literature, eg Campbell and Mankiw (1989, 1991).

This implies the following aggregate consumption function:

$$C_t = b\left(HW_t^{PIH} + FW_t\right) + YL_t^{C-C} \qquad \textbf{(9)}$$

where HW_t^{PIH} is total human wealth of PIH consumers, FW_t is total net non-human wealth and YL_t^{C-C} is total labour income of the credit-constrained. We assume here that the real interest rate used by PIH consumers to discount their expected future labour income is unchanged when there are credit-constrained individuals in the economy. We further assume that the trend growth of consumption (the drift term, g) is also independent of whether or not there are credit-constrained consumers in the economy. These are, of course, simplifying assumptions. For example, the existence of credit constraints for some individuals may negatively affect overall productivity growth and hence g, though arguably there might only be a levels effect. More importantly, b, HW_t^{PIH} and FW_t are all functions of the real interest rate which is itself likely to be affected by capital market imperfections. But the aim here is to investigate the consequences of adopting a simple treatment of credit-constrained consumers as in Campbell and Mankiw (1989, 1991), where differences in b and g are not explicitly allowed for. And although the strength of the result in this section is sensitive to these assumptions, the main conclusion (that simplistic treatments of credit constraints cannot necessarily explain the excess sensitivity of consumption to human wealth) holds even when these assumptions are relaxed.

It is also commonly assumed for simplicity that PIH consumers receive a constant proportion, l, of total labour income. This is equivalent to saying that the growth rate of labour income for credit-constrained and unconstrained consumers is equal. But there are good reasons why in practice credit-constrained consumers might actually have above-average expected labour income growth, eg if the constrained are young individuals who have accumulated no financial savings and so would like to borrow off their (high) expected labour income growth. I investigate the consequences of this simplifying assumption by further assuming that human wealth can be approximated by a linear function of labour income (with no additional terms), as in equation **(7)**. In this case, the aggregate consumption function can be written as:

$$C_t = b\left(l YL_t \left(\frac{1+r}{r-g} \right) + FW_t \right) + (1-l) YL_t$$

or

$$C_t = bFW_t + \left(b\left(\frac{1+r}{r-g} \right) l + (1-l) \right) YL_t \qquad \textbf{(9')}$$

where YL_t is aggregate labour income. This can be compared with the aggregate consumption function when all individuals are PIH consumers:

$$C_t = b\left(\left(\frac{1+r}{r-g} \right) YL_t + FW_t \right) \qquad \textbf{(10)}$$

It is now straightforward to compute the elasticity of consumption with respect to labour income in the two different economies using equations **(9')** and **(10)**:

$$\hat{b} = \frac{\P C_t}{\P YL_t} . \frac{YL_t}{C_t} = \left(b\left(\frac{1+r}{r-g} \right) l + (1-l) \right) \frac{YL_t}{C_{tt}} \qquad \textbf{(11)}$$

$$b = \frac{¶C_t}{¶YL_t} . \frac{YL_t}{C_t} = b\left(\frac{1+r}{r-g}\right)\frac{YL_t}{C_t} \qquad \textbf{(12)}$$

Now consider two mutually exclusive hypotheses:

The first hypothesis is that the PIH is consistent with the value of the econometrically-estimated elasticity, and that there are no credit-constrained individuals. We can substitute in values for the real interest rate, the annuity rate, the drift term g, aggregate labour income and consumption into **(12)**, and compare this with the econometric estimate. But we have shown earlier that **(12)** is equivalent to the human wealth share under the PIH, and have already concluded that calibrated estimates of this share are too low to explain the econometric estimates of the elasticity. So we reject this hypothesis.

The second hypothesis is that the econometric estimate of the elasticity can be reconciled with theory if in fact some individuals are credit-constrained, and so the pure PIH is incorrect.(23) We can see whether assuming some individuals are credit-constrained will help to bridge the gap between the PIH and the econometric estimate by comparing the expressions in **(11)** and **(12)**. This gives a condition for when the elasticity of consumption with respect to labour income is higher in the model where there are credit-constrained consumers:

$$(13) b\left(\frac{1+r}{r-g}\right) < 1, \, where b \approx r - \frac{(r-d)}{q}\frac{1}{(1+r)(1+d)}$$

(23) Of course, I am implicitly assuming in this that the modeller cannot observe the financial status of individuals separately - otherwise he could see directly whether any individuals are credit-constrained.

But equations **(2)** and **(7)** show that $b\left(\frac{1+r}{r-g}\right)$ is actually the marginal propensity to consume out of labour income for PIH consumers. So condition **(13)** states that the elasticity of consumption with respect to labour income when some individuals are credit-constrained consumers is higher than if all individuals are PIH consumers if the marginal propensity to consume out of labour income for PIH consumers is less than one.

This result is consistent with intuition: equation **(8)** says that the marginal propensity to consume out of labour income for credit-constrained consumers is unity. So if some individuals in the economy are in fact credit-constrained, other things being equal, this implies a higher estimate of the marginal propensity to consume at the macro level.

However, condition **(13)** does *not* in general hold for plausible values of the growth rate, return on savings, discount rate and intertemporal elasticity of savings. For example, if $r=\delta$ this condition *never* holds: the marginal propensity to consume out of labour income for PIH consumers is always greater than or equal to one. Intuitively, if labour income follows a random walk with drift, all shocks to the level of income in the long run are permanent. It is perhaps not surprising in these circumstances that shocks to current income lead to an upward revision of permanent income by PIH consumers so that the marginal propensity to consume is in general greater than one. This means that in this simple model, an economy with credit-constrained consumers will in general have a lower marginal propensity to consume and a lower elasticity, than if all individuals are PIH consumers.

This result is striking and may seem counterintuitive. The elasticity of consumption with respect to labour income for PIH consumers, if log labour income follows a random walk with drift, is equal to the share of human wealth in total wealth, say 0.78. The elasticity for credit-constrained consumers is unity. What we have shown is that if the economy is in fact populated by a combination of PIH and credit-constrained consumers, then the aggregate elasticity would not in general lie between 0.78 and unity!

Such a result might not seem possible in a linear model. The explanation is clearer when it is recognised that if the economy were populated by some credit-constrained individuals, and such individuals are assumed by definition to hold no net non-human wealth, then the per capita non-human wealth holding of the PIH consumers would be greater than in the case where all individuals are PIH consumers. Other things being equal, this implies that the human wealth share for PIH consumers is lower than it otherwise would be if all individuals were unconstrained. So the aggregate elasticity would not be a linear combination of unity and 0.78, but of unity and some number lower than 0.78. One should avoid the temptation of incorrectly treating the elasticity for PIH consumers as a parameter when it is in fact endogenous.

What this means is that the simple treatment of exogenous credit constraints above is not able to explain the excess sensitivity of consumption to innovations in labour income detected in Section 3. Macro modellers will require a more sophisticated treatment of credit constraints, where the constraints are endogenous, if they are to explain the excess sensitivity.

5 Conclusion

Under certain assumptions the permanent income hypothesis (PIH) argues that individuals base their consumption decisions on lifetime wealth, not current income. The levels consumption function generates at least two testable predictions. First, that a regression of log consumption on log total wealth, controlling for the annuity rate, should give a coefficient of unity.

Second, that the elasticity of consumption with respect to human wealth should equal the share of human wealth in total wealth.

This paper has tested the second prediction. Assuming that individuals expect their labour income to follow a random walk with drift (a model that seems to fit the aggregate data well) the PIH predicts that the elasticity of consumption with respect to labour income is equal to the share of human wealth in total wealth.

Our baseline estimates suggest that the human wealth share has averaged 0.78 over the period 1975 Q1-1998 Q4. And if risk-averse agents discount their future labour income at a higher rate, as in Blake, Camba-Mendez and Weale (1998), then this 0.78 figure should probably be viewed as an upper bound. Elasticities with respect to labour income in aggregate consumption functions estimated over broadly the same sample period are (significantly) higher than this: consumption appears to be more sensitive to changes in labour income than is predicted by the PIH. The existence of credit constraints is an obvious candidate to explain the discrepancy between the PIH and the empirical evidence. But a simplistic treatment of credit constraints sometimes used in the literature, assuming that there are credit-constrained individuals (who cannot borrow and do not save) with a constant proportional claim on national labour income, does not explain the difference. There are more rigorous attempts at deriving solved-out consumption functions when there are credit constraints in the literature (see eg King (1986)), and these are likely to form more appropriate justifications for the excess sensitivity of consumption detected in consumption functions estimated on aggregate data.

Annex

The optimisation problem for the infinitely-lived representative agent in a world of no uncertainty is to maximise discounted lifetime utility subject to a lifetime budget constraint:

$$\underset{c_t}{Max}\sum_{s=t}^{\infty}\frac{U(c_s)}{(1+d)^{s-t}} \quad \text{s.t.} \quad \sum_{s=t}^{\infty}\frac{c_s}{(1+r)^{s-t}}=\sum_{s=t}^{\infty}\frac{yl_s}{(1+r)^{s-t}}+fw_t\,(=w_t)\,,$$

where c_s is consumption of goods and services in period s (consumption of housing services are considered in the empirical section), yl_s is labour income in time s, fw_t is financial wealth in period t, w_t is total wealth in period t, d is the agent's subjective rate of time discount and r is the real return to savings (assumed constant for simplicity).

The first-order conditions for this problem give the following expression relating the growth in the marginal utility of consumption to the difference between the discount rate and the return to savings (the Euler equation):

$$\frac{U'(c_t)}{U'(c_s)}=\frac{(1+r)^{s-t}}{(1+d)^{s-t}}$$

Assuming CRRA preferences, $U(c_t)=\frac{c_t^{1-q}}{1-q}$, where $\frac{1}{q}$ is the intertemporal elasticity of substitution and $q>0$, the Euler equation becomes:

$$\left(\frac{c_s}{c_t}\right)^q=\left(\frac{1+r}{1+d}\right)^{s-t}$$

Substituting this into the lifetime budget constraint gives:

$$\sum_{s=t}^{\infty} c_t \left(\frac{1+r}{1+d}\right)^{\frac{s-t}{q}} \cdot \frac{1}{(1+r)^{s-t}} = w_t$$

$$\therefore c_t \sum_{s=t}^{\infty} \left(\frac{1+r}{1+d}\right)^{\frac{s-t}{q}} \cdot \frac{1}{(1+r)^{s-t}} = w_t$$

$$\therefore c_t \left\{ 1 + \left(\frac{1+r}{1+d}\right)^{\frac{1}{q}} \cdot \frac{1}{1+r} + \left(\frac{1+r}{1+d}\right)^{\frac{2}{q}} \cdot \frac{1}{(1+r)^2} + \dots \right\} = w_t$$

$$\therefore c_t = b \, . w_t \, ,$$

where $b = 1 - \frac{1}{1+r} \cdot \left(\frac{1+r}{1+d}\right)^{\frac{1}{q}}$. Taking a first-order Taylor approximation of $\left(\frac{1+r}{1+d}\right)^{\frac{1}{q}}$ around $r = d$, it is easy to show that $b \approx \frac{r}{1+r} - \frac{(r-d)}{q} \frac{1}{(1+r)(1+d)} \approx r - \frac{(r-d)}{q} \frac{1}{(1+r)(1+d)}$ for small values of r. The first term can be interpreted as the income effect from a change in r while the second term captures the substitution effect.

Total wealth is comprised of discounted future labour income (human wealth, hw_t) and discounted future property income (non-human wealth, fw_t.). By substitution this gives the following equation for the consumption function:

$$c_t = b\left(hw_t + fw_t\right)$$

References

Bank of England (1999), *Economic Models at the Bank of England.*

Blake, A P, Camba-Mendez, G and Weale, M (1998), 'UK consumption in the long run: determinants of consumer spending 1925-1995', *National Institute of Economic and Social Research*, Discussion paper No 138.

Blanchard, O J (1985), 'Debts, deficits and finite horizons', *Journal of Political Economy*, 93, pages 223-47.

Blanchard, O and Fischer, S (1989), *Lectures in Macroeconomics*, MIT Press.

Buiter, W H (1988), 'Death, birth, productivity growth and debt neutrality', *Economic Journal*, 98, pages 279-93.

Campbell, J Y and Mankiw, N G (1989), 'Consumption, income and interest rates: reinterpreting the time-series evidence', *NBER Macroeconomic Annual*, pages 185-215.

Campbell, J Y and Mankiw, N G (1991), 'The response of consumption to income: a cross-country investigation', *European Economic Review*, 35, pages 723-767.

Chan, A, Savage, D and Whittaker, R (1995), 'The New Treasury Model', *Government Economic Service Working Paper No 128*.

Darby, J and Ireland, J (1994),'Consumption, forward-looking behaviour and financial deregulation', *University of Strathclyde Discussion Paper No 20*.

Deaton, A (1992), *Understanding Consumption*, Clarendon Press.

Ellwood, K S (1998), 'Testing for excess sensitivity in consumption: a state-space/unobserved components approach', *Journal of Money, Credit and Banking*, 30, pages 64-82.

Flavin, M (1981), 'The adjustment of consumption to changing expectations about future income', *Journal of Political Economy*, 89, pages 974-1,009.

Hall, R E (1978), 'Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence', *Journal of Political Economy*, 86, pages 971-88.

Hayashi, F (1982), 'The Permanent Income Hypothesis: estimation and testing by instrumental variables', *Journal of Political Economy,* Vol 90, pages 895-916.

King, M (1986),'Capital market 'imperfections' and the consumption function', *Scandinavian Journal of Economics*, 88(1), pages 59-80.

Macklem, R T (1997), 'Aggregate wealth in Canada', *Canadian Journal of Economics,* Vol 30.

Muellbauer, J (1994), 'Consumption', *Oxford Review of Economic Policy*.

Muellbauer, J (1996), 'Income persistence and macro-policy feedbacks in the United States', *Oxford Bulletin of Economics and Statistics*, 58:4, pages 703-33.

Muellbauer, J and Lattimore, R (1995), 'The consumption function: a theoretical and empirical overview' in Pesaran, M H and Wickens, M (eds), *Handbook of Applied Econometrics*, Blackwell, Oxford.

National Institute for Economic and Social Research (1999), *Domestic model manual*.

Price, S (1999), 'Risk-adjusted human capital and UK consumers' expenditure', *City University Discussion Paper*.

Sheldon, M and Young, G (1999), 'Explaining aggregate consumers' expenditure in a changing financial environment: A comparison and synthesis of different approaches', *unpublished mimeo*.