

A RETROSPECTIVE ON FRIEDMAN'S THEORY OF PERMANENT INCOME*

Costas Meghir

Friedman's book on the consumption function is one of the great works of Economics demonstrating how the interplay between theoretical ideas and data analysis can lead to major policy implications. We present a short review of Friedman's permanent income hypothesis, the origins of the idea and its theoretical foundations. We give a brief overview of its influence in modern economics and discuss some relevant empirical results and the way they relate to the original approach taken by Friedman.

Friedman's (1957) book on the consumption function is one of the great works of economics demonstrating how the interplay between theoretical ideas and data analysis could lead to major policy implications. The theory of the consumption function played an important role in explaining why traditional Keynesian demand management, through transitory tax policy or other transitory income boosting measures can have little or no effect on real consumption and on the desired policy outcomes. The apparently simple ideas in this excellent book have been so insightful and powerful that they have given rise to a huge amount of research, both theoretical and empirical, which continues to this date. In this short article I trace out some of the research relating to the permanent income hypothesis (PIH), particularly that which has been based on microeconomic data, and I demonstrate the relevance of these ideas for our current way of thinking about consumption, savings and income processes.

Friedman and Kuznets (1954), which was actually written in the early 1940s but delayed in publication, first formulated many of the ideas of the PIH, including the permanent/transitory decomposition of income. The core of these ideas, together with further tests, was brought together in Friedman (1957). In my review of the PIH I draw mainly from this latter work.

1. The Permanent Income Hypothesis

1.1. A Statement of the Hypothesis

Milton Friedman's PIH originates from the basic intuition that individuals would wish to smooth consumption and not let it fluctuate with short-run fluctuations in income. In fact the model was developed to explain important empirical facts in a unified framework. For example, why is income more volatile than consumption and why is the long-run marginal propensity to consume out of income higher

* This article was prepared for a conference in honour of Milton Friedman at the University of Chicago in November 2002, on the occasion of his 90th Birthday. I thank the organisers for the invitation. I also thank Orazio Attanasio, Richard Blundell, Martin Browning, Jim Heckman, Hide Ichimura, an anonymous referee and colleagues at IFS and UCL for useful discussions that helped me organise this presentation. I am of course responsible for all errors and interpretations.

than the short-run one? To answer these questions Friedman hypothesised that individuals base their consumption on a longer-term view of an income measure, perhaps a notion of lifetime wealth or a notion of wealth over a reasonably long horizon. The basic hypothesis posited is that individuals consume a fraction of this *permanent* income in each period and thus the average propensity to consume would equal the marginal propensity to consume. The propensity itself could vary with a number of factors, including the interest rate and taste shifter variables, or could reflect uncertainty we will return to these important insights below.

Friedman set himself the task of testing his hypothesis against an increasing set of empirical facts from time series data and budget studies. The standard least squares regression of consumption on income would always point to a marginal propensity to consume below the average propensity. Conditioning on extra regressors seems to make things worse. It is at this point that Friedman's ingenuity, brought together the literature on budget studies by Reid (1952), Morgan (1951) and others, as well as time series analyses with econometric ideas on measurement error to devise estimation techniques, that not only allowed the testing of the basic hypothesis, but led to the estimation of underlying parameters that directly characterised the Permanent Income Hypothesis (PIH). As far as the role of measurement error is concerned the works of great economists and statisticians of the time, namely Harold Hotelling¹ and James Durbin, influenced Friedman. He brings together empirical results and statistical theory, and develops new results by combining these with his economic ideas.

The ingredients of Friedman's model are permanent consumption (c_p), permanent income (y_p), transitory consumption (c_t), transitory income (y_t). Measured income is the sum of *permanent* and *transitory* income (y_t) and measured consumption is the sum of *permanent* and *transitory* consumption (c_t), i.e.

$$c = c_p + c_t$$

and

$$y = y_p + y_t.$$

Permanent consumption is determined by the equation

$$c_p = k(r, \mathbf{z})y_p$$

where $k(r, \mathbf{z})$ is the average (or marginal) propensity to consume out of permanent income which depends on the rate of interest and on taste shifter variables \mathbf{z} . The transitory components may reflect genuine fluctuations, or measurement errors. The key point is that the consumption plan does not depend on the transitory components. To provide empirical content to this hypothesis, Friedman added the assumptions that the transitory components are uncorrelated to each other and uncorrelated to the permanent component.

Friedman then goes on to show that the slope coefficient of a regression of observed consumption on observed income would lead to an underestimate of the marginal propensity and to a positive estimated intercept. Moreover he shows that

¹ See Hotelling (1933) and Friedman (1992) on the regression fallacy.

the rate of attenuation of the marginal propensity to consume is equal to the ratio of the variance of the permanent income to total income in the data at hand (cross section or time series). This is now known as attenuation bias due to classical measurement error. He also showed that the rate at which k is attenuated is equal to the elasticity of observed consumption to observed income, which provided him with an independent test of the PIH, since he could estimate the same parameter by two different ways, using different data sets.

Using this simple but ingenious apparatus, he set out to explain the results from budget studies and time series data that seemed to contradict the permanent income hypothesis.

Why is it the case that the marginal propensity to consume in cross sectional data is lower than the average propensity and how can this be reconciled with the fact that the consumption function shifts upwards over time? His interpretation was that the results from a single cross section suffer from attenuation bias and the estimated consumption function using observed data shifts over time because permanent income goes up. Thus for a given level of observed income, permanent income is higher in later years than in earlier ones. In fact if one joins the average observed pairs of consumption-income points across time one recovers a function that implies that the marginal propensity is equal to the average one, in the data that Friedman used. The idea is that average income reflects (average) permanent income, since by the law of large numbers the transitory components average out. The pattern is repeated in a number of other data sets and circumstances. For example an interpretation of why blacks save more than whites with the same observed income is that the former have lower permanent income than whites. When we condition on blacks being in the same *observed* income class, the transitory component of income (which does not affect consumption) is likely to be larger among blacks than whites. Similar arguments can be made when we compare the self-employed to the salaried workers or farm to non-farm households, the first in each pair having larger transitory components to their income.

The evidence he presented in favour of the PIH is compelling; however it is interesting to interpret the method of analysis that Friedman used with modern terminology and to state explicitly the assumptions made, since this will allow us to motivate the subsequent research on consumption.

By comparing averages over time or across different groups of individuals, Friedman is implicitly using instrumental variables, now a well recognised technique for dealing with *classical* measurement error in linear models.² Take the comparisons over time. The basic premise is that because of overall growth, later time periods are associated with higher permanent income. Moreover, if neither the mean of the transitory component of income nor preferences change over time we can use time as an instrument; the procedure used is what has recently been used to describe the generalised Wald estimator. Wald (1940) suggested comparing group averages to overcome measurement error, although the way he

² See Goldberger (1972) for an account of instrumental variables as a method for dealing with endogeneity and errors in variables with non-experimental data. Angrist and Krueger (2001) give a recent account of the use of the technique.

approached grouping was not quite right. This is because he defined groups by class intervals of the variable that was supposed to be measured with error. The group averages will not be devoid of measurement error in this case. Friedman's approach, in the same spirit, was actually correct and bypassed Wald's error because he constructed groups based on other variables, such as time or ethnicity or occupation. Moreover, what Friedman did informally, joining means across groups, Eisner (1958) implemented formally by using as comparison groups occupation, city class categories and/or age either alone or crossed. When regressing group means of consumption on group means of income Eisner (1958) found a remarkable agreement of the facts with the PIH, as all elasticities in the between group regressions were very close to one.

There may be reasons to question some of the assumptions underlying the empirical analysis of both Friedman and Eisner, as is often the case. For example the analysis relies on whether we believe preferences (and other factors determining consumption) are similar enough across the groups we are comparing or across time. Nevertheless, this is one of the first theoretically coherent applications of instrumental variables to empirical analysis, with extraordinarily successful results.

2. The Notions of Permanent Income

The PIH provides a flexible framework for the study of consumption and savings. Much of the subsequent research has filled in those elements necessary for explaining aspects of the data, thus defining the agenda for research on consumption and savings. Contrary to formal versions of the model, which I discuss later, the PIH as stated by Friedman did not take a very firm view on the appropriate horizon for consumer choice. The basic hypothesis requires a rolling horizon, which allows some degree of smoothing of consumption. However, the flexibility of the hypothesis comes at a price, since it is hard to define the theoretical underpinnings of the model and consequently it is hard to make more detailed statements about policy.

The loose definition of permanent income leaves open the question of its measurement. A moving average or distributed lag of past incomes are both options that have been suggested. However the way one measures permanent income should be part of the model and has to do with the way individuals form expectations as well as with the stochastic process underlying income. Friedman points this out in the consumption function book when he states that that permanent income is best defined '... to be whatever seems to correspond to consumer behaviour.' (p.23) This invites us to use theory to identify from the data on consumption the correct notion of permanent income, or as happened eventually, to bypass its measurement completely, when testing the hypotheses corresponding to some version of the PIH. The development of the Euler equation approach for the study of savings and consumption, the introduction of Rational expectations in economic modelling and the Lucas (1976) critique, provided both the motivation and the means by which to develop tests of the PIH that did not involve measuring PI, at least with backward looking distributed lags, or other similar tools.

We will now turn to describing the theoretical foundations of the hypothesis and the empirical analysis that followed.

3. The Theoretical Foundations of the PIH

There have been other attempts to explain consumption behaviour. The reason the PIH has endured so much is because, beyond its simple intuitive appeal, it is a special case of an intertemporal optimisation model of consumer behaviour, which is the most coherent and logically consistent model we have at present.³ This model has its roots in the works of Fisher (1907) and Ramsey (1928) and has since been developed in many directions.

Let us recall the basic elements of this model. Since the PIH includes the notion of shocks to income, which may reflect real but transient or even permanent fluctuations to income, the theoretical framework should allow for uncertainty.

Consider a consumer who maximises expected lifetime utility, subject to a life-time budget constraint. Expectations are rational and there are no liquidity constraints, in the sense that the individual can borrow and lend at a constant interest rate. Thus we have that the optimal consumption plan at time zero maximises the sum of expected utility subject to a life-time budget constraint.

$$\max_{c_t, t = 1, \dots, T} \left[\mathbf{E}_0 \sum_{t=1}^T \beta^t U(c_t) \right]$$

subject to

$$A_0 + \sum_{t=1}^T R^t (y_t - c_t) = 0$$

where A_0 is initial non-human wealth. Under uncertainty the problem is solved in each period to accommodate news on income. The key implication of this model

$$U'(c_t) = \frac{R}{\beta} \mathbf{E}_t \lambda_t$$

is that the marginal utility of consumption in each period is equal to the expected marginal utility of wealth multiplied by a factor depending on the interest rate and the rate of time preference. This expresses very closely the ideas underlying the PIH. Under certain circumstances this yields a version of the PIH, where permanent income is defined as total lifetime wealth.

Thus for example, if the rate of return is equal to the discount rate ($R = \beta$) and there is no uncertainty, consumption will be constant if preferences do not change over time. In this case we can write consumption as a constant proportion to total wealth, which provides a version of the PIH when transitory shocks are all meas-

³ There have been many models of consumption. An influential example is the Keynesian consumption function (Keynes, 1935) based as Keynes put it on a 'basic psychological law'. However his consumption function lacks any microeconomic foundations based on individual optimisation. An alternative is Duesenberry (1949) which was based on the idea that relative positions in the income distribution matter for utility.

urement error. However a better insight can be obtained by using the quadratic utility function. With quadratic utility and assuming expected utility maximisation we get that consumption is expected to remain the same

$$c_t = E_t c_{t+1}$$

which formed the basis of a test of the PIH by Hall (1978). This combined with the budget constraint gives the result that consumption is proportional to expected wealth

$$c_t = \frac{r}{[(1+r) - (1+r)^{-T}]} E_t W$$

where $E_t W = A_t + E_t \sum_{s=t}^T R^s y_s$ is the expected life-cycle wealth. This provides an immediate justification for the PIH and shows how the average (marginal) propensity to consume relates to the interest rate. It is also very closely related to the ideas expressed in the work of Modigliani and Brumberg (1954) and Ando and Modigliani (1963) on the Life-Cycle Hypothesis, which introduces formally life-cycle considerations.

Now lets see how consumption is supposed to react to changes in current income, which was one of the key motivations for the PIH. As shown by Sargent (1978), Flavin (1981) and later elaborated by Campbell (1987) this version of the PIH implies that consumption changes are equal to the annuity value of all revised changes in future incomes. Thus for the infinite horizon case we can write

$$\Delta c_t = \frac{r}{1+r} \sum_{s=0}^{\infty} \frac{1}{(1+r)^s} (E_t - E_{t-1}) y_{t+s}$$

The term $(E_t - E_{t-1}) y_{t+s}$ reflects revisions in expectations on the income flow. In the simplest case where the income process contains a deterministic component possibly time varying but known in advance, plus a transitory component we get that consumption does not react (much) to current income fluctuations. In fact the reaction is zero for anticipated changes and the reaction to unanticipated transitory shocks equals the annuity value of the shock. This will typically be very small, at least for a young individual. This reflects very clearly the thinking of Friedman and it is an intuition that finds expression in modern economic analysis. However, there are doubts on whether this is a good description of the income process. If the shocks to income are permanent, then all future levels of income will be revised upwards or downwards by the same amount leading to a reaction of consumption, which is equal to the change in current income. This is perfectly consistent with the PIH, since Friedman explicitly argued that changes over time should reflect (unexpected) changes in permanent income.

In discussing permanent income Friedman explicitly allowed for the possibility for some form of Liquidity constraints, when he stated that the discounted sum of future incomes does not have much intuitive appeal as a measure of permanent income because it '... implies that units can borrow [...] at the same interest rates at which they can lend ...' (p.25). So how do liquidity constraints affect the argument and are they really important from an empirical point of view?

With liquidity constraints in the model discussed above, individuals with positive assets will always behave as in the PIH but individuals with zero assets may have consumption behaviour such that consumption tracks predictable changes in income, whether transitory or not, in the presence of borrowing restrictions. Whether this will be detectable in the data will depend on the proportion of individuals who are thus constrained.

Thus on closer examination it follows that the PIH and its implications for policy requires us to take a stance on the definition of permanent income and on how capital markets operate.

These observations have given rise to a large literature on testing for LCs based both on macroeconomic aggregate data and on microeconomic data. Two of the best known tests of LCs on aggregate data are the one by Flavin (1981) and Blinder and Deaton (1985). Flavin (1981) estimates a structural time-series model of consumption and income and tests whether lagged income growth matters for consumption, once one controls for innovations to consumption due to unexpected changes or 'surprises' in current income. Blinder and Deaton (1985) implement a similar 'excess sensitivity' test which states that under the null only income surprises should matter for consumption growth. Predictable income growth should not matter. The findings of both these papers is that the version of the PIH with perfect capital markets, as expressed in the equation above, is rejected – in line with Friedman's suspicions. In both papers, predictable changes in income are shown to affect consumption. Nevertheless, the sensitivity of consumption to these predictable income changes is actually quite low, with a marginal propensity to consume out of current income in the Blinder and Deaton paper of about 0.15, once we control for PI. This is not too say that this is not important but to lay the ground for arguing that a more general version of the PIH, with similar foundations, is in fact consistent with the data, although the version of the PIH developed here is probably too restrictive.

4. Risk Aversion, and the PIH

A key restriction of the model above is the absence of any considerations of risk aversion. Given that markets are never found to be complete (Cochrane, 1991; Attanasio and Davis, 1996) and given aggregate uninsurable shocks this must be a central issue.

Friedman was explicit about the role both of taste shifter variables, such as demographic composition and about uncertainty, in determining the average/marginal propensity to consume. He saw that under uncertainty, the average propensity to consume out of permanent income could be lower. He also saw that the average propensity to consume may differ depending on demographic composition of the household. While these ideas were not formalised at the time they turn out to be key elements in demonstrating that some generalised version of the PIH does in fact fit the data very well and can explain observed consumption patterns. Friedman however, rightly warned against 'controlling' in a simple minded way for demographic effects or other variables that may proxy for permanent income, since these

could exacerbate the impact of measurement error and falsely imply that the PIH does not work. This ‘partial correlation analysis’, as he calls it, can in his words render a budget study ‘worthless’ however much ‘loving care’ is put into the analysis. This issue recurred in the work of Griliches (1977) many years later in the context of studying the impact of education on wages, when education has been measured with error. The argument indicates that allowing for taste shifter variables when testing the PIH is a delicate matter and may undermine the hypothesis. So we will now describe an extended model applied to microeconomic data and we will discuss the results obtained in this framework.

Consider a framework, where preferences for consumption depend on household composition. It may for example be reasonable to think that more consumption is shifted towards the household when it consists of more members. Another important mechanism for this is the price index itself; if households with children consume a different bundle of commodities than households without then the relevant deflator would be different.

The quadratic preference specification, used to show the theoretical foundations of the pure version of the PIH, does not allow for precautionary savings. The individual behaves in exactly the same way as if she were given with certainty the expected lifetime wealth. This is of course very restrictive. So, in addition to controlling for changes in preferences as a function of household circumstances we also change the structure of the utility function to have a positive third derivative, which has been proven to be the key to allowing for precautionary savings (Sandmo, 1969; Drèze and Modigliani, 1972).

Most of the ideas can be presented in a framework where the utility function takes the constant relative risk aversion form with a discount rate or a risk aversion coefficient which is a function of demographic variables and written as

$$U(c_t|z_t) = \frac{1}{\rho + 1} a(z_t) c_t^{\rho+1} \quad \rho < 0$$

where $a(z_t)$ is a function of variables that can affect preferences. The choice of these variables is of course a key issue. First, one might include demographic variables, such as family size and the number and ages of children, since these will shift the needs and preferences of the household over time. Zeldes (1989) bases a test for liquidity constraints on a model including variables that reflect needs and rejects the model. However, his model is based on food consumption alone. This specification is likely to reject the hypothesis of no liquidity constraints, unless preferences across commodities are additively separable and homothetic. This is because we will be attributing changes in consumption growth that are due to changes in prices, to changes in real incomes. Attanasio and Weber (1995) show the practical importance of this issue.

However, this is not the only reason why the model was rejected. Heckman (1974) pointed out that if consumption is not additively separable from hours of work and this is ignored, the resulting consumption choices over time will look as if they track income, in a way that would reject the basic premise of the permanent income hypothesis. This suggests that we should include variables describing the labour supply behaviour of the household members.

It is not in general possible to find an explicit solution to the consumption function when preferences are isoelastic, as in the example above. Moreover it is probably not desirable for the purposes of testing the model, since this would require a complete specification of the income process, information about assets and a modelling of the evolution of demographics. Thus the empirical approach to test for this model is to use the Euler equation. This can be approximated by

$$\Delta \ln c_t = \theta r_t + \theta(E_{t-1}\sigma_t^2 - \delta) + \beta' \Delta \mathbf{z}_t + u_t$$

This expresses most of the intuitions in Friedman's PIH *and* as we shall argue, fits the facts very well. According to this equation consumption growth increases when interest rates are predicted to be high. Moreover consumption growth will be positive, implying higher savings, when the conditional variance of consumption is higher than the personal rate of discount (the trade-off between precaution and impatience). The other source of changes in consumption growth, are changes in taste shifter variables and labour supply, summarised in \mathbf{z} . Finally the error term u_t reflects changes due to unexpected events. As before, these will be equal to the permanent shocks to income and will contain a very small fraction of a transitory shock, for a reasonably young person.

This formulation gets round the problems of the 'partial correlation analysis' discussed by Friedman because we implicitly control for permanent income. This regression can be seen to be the log difference of Friedman's basic equation, where the permanent income (or in terms of a more recent terminology the inverse of the marginal utility of wealth) is differenced out and hence does not need to be measured.

We use the graphs in Figure 1 (taken from Blundell *et al.*, 1994) to show intuitively why allowing for changes in preferences over the lifecycle is key to showing that the PIH is a very good description of the data. In these graphs we show log consumption, log income, the number of children and female participation, over the life-cycle of individuals. It has been constructed by overlapping cohorts over a long period of time, in the spirit of Ghez and Becker (1975). Each line on a graph represents the time series evolution of consumption, income, demographics and female participation for one cohort. This gives us a synthetic view of how key variables change over the life-cycle.

The first interesting point is that income appears to be more volatile than consumption, since the lines for log income are more dispersed than those for consumption, implying that individuals smooth out fluctuations that are not perceived as permanent. Moreover, the large income growth apparent for the middle cohorts, is not matched by equivalent income growth for consumption, also implying that this growth was perceived as transitory. However, the observation of those suggesting that the PIH is invalid is that the graph of household consumption looks remarkably like the graph of household income, despite the fact that the income path is predictable and hence does not represent changes in permanent income. Interestingly, when female participation dips in the child-bearing ages, consumption does not dip. However, neither does income, telling us something about the timing of births relative to male earnings growth but little

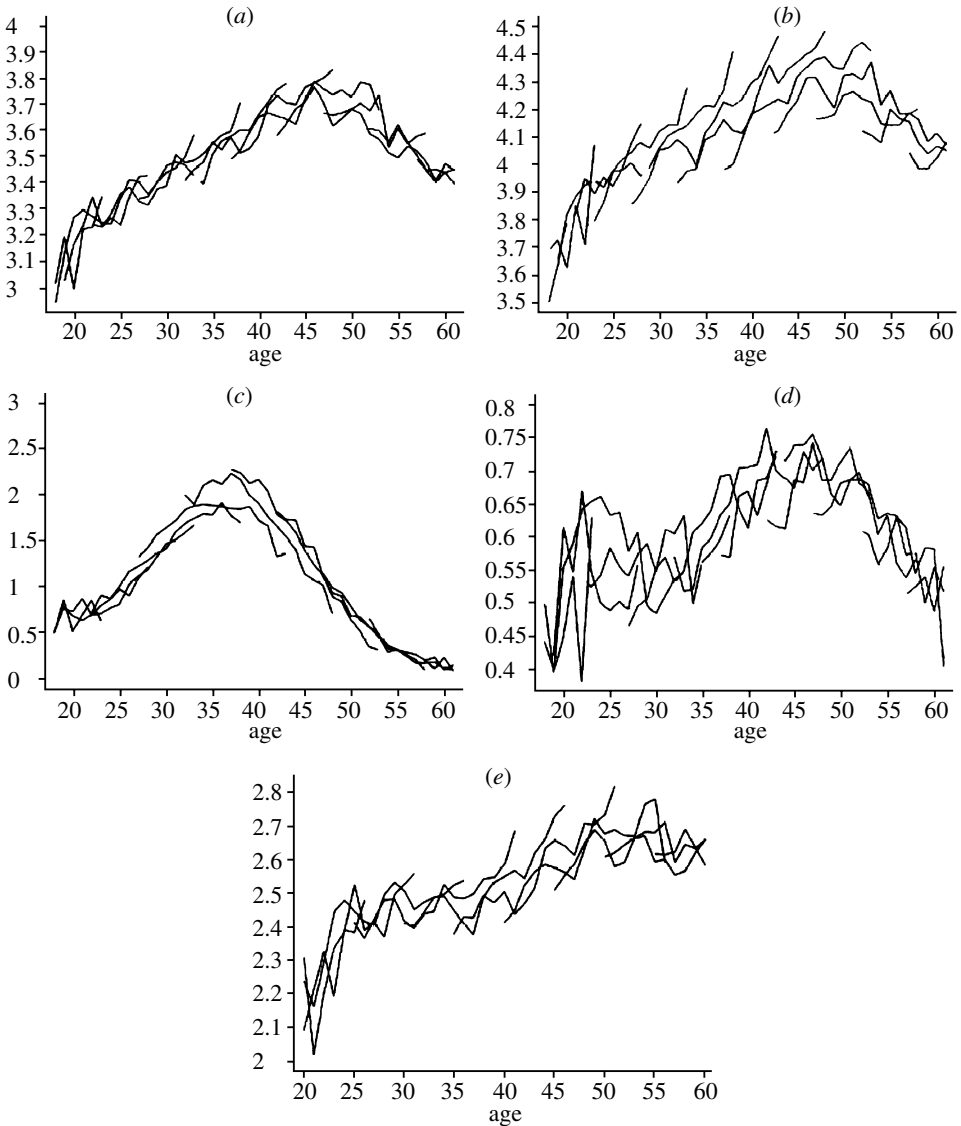


Fig. 1. (a) Log real household consumption over the life cycle. (b) Log real household income over the life cycle. (c) Number of children in the household over the life cycle. (d) Female employment over the life cycle. (e) Log consumption per equivalent adult over the life cycle. Source: Blundell *et al.* (1994)

about the validity of the PIH. However, we also note that the number of children also tracks income. Indeed log consumption *per adult* equivalent allowing for the change in household size, grows more or less steadily, which is consistent with the presence of precautionary savings and with the PIH, as I will illustrate below.

The equivalence scale is of course arbitrary and does not allow for the all-important impact of labour supply. We can include these variables freely and we

can allow for the fact that they are endogenously chosen, when estimating the model for consumption. The question one asks empirically is whether the estimated model fits the data and whether the overidentifying restrictions are acceptable. So in Table 1 we show the impact of predictable income growth on consumption growth as estimated in Blundell *et al.* (1994) for the UK and Attanasio and Weber (1995) for the US. In both cases the effect of predictable income growth on consumption growth declines substantially and becomes insignificant when we allow for demographics and labour supply. Thus most of the 'income tracking' behaviour of consumption can be explained away when we allow consumption to be non-separable from labour supply. In the Blundell *et al.* (1994) paper, it seems that most of the action comes from controlling for male labour supply. Importantly, in both papers, including these other factors, from which consumption is non-separable, reduces the coefficient of income growth very substantially, as well as making it statistically insignificant. In the Blundell *et al.* (1994) paper, this occurs without a substantial loss in precision.

From this, we can conclude that once we control for uncertainty and for non-separability this extended PIH describes the data very well. Nevertheless, one issue that may put the result into question is whether the test still has power, because transitions in and out of employment are an important source of income fluctuations. On the other hand (predictable) wage rate fluctuations are also important and provide information for testing for the PIH. This is particularly true over recent periods with such large changes in education and gender wage differentials, some of which seem to be predictable.

5. The Nature of Permanent Income and the Estimation of the Time Series Properties of Income

Quite apart from the study of consumption, Friedman's work also set the agenda for studying income processes. The study of income processes has been motivated in part by the study of consumption and savings because it has become apparent that knowing the time series properties of income may be informative about consumption behaviour. It can for instance possibly provide additional testing power for economic hypotheses. Friedman demonstrated this when he used the income process to obtain an indirect estimate of the elasticity of consumption to current income and thus provide an additional test of his hypothesis.

Table 1
Excess Sensitivity Tests in the UK and the US using Micro Data
(Dependent variable: Consumption growth)

	No labour supply No demographics	Demographics only	Labour supply and demographics
UK: Blundell <i>et al.</i> (1994) Predicted income growth	0.537 (0.196)		0.143 (0.142)
US: Attanasio and Weber (1995) Predicted income growth		0.247 (0.058)	0.100 (0.103)

Lillard and Willis (1978) were one of the first to study income processes. They modelled income as the sum of a deterministic part, including a trend and other variables reflecting human capital and of an unobserved component including a fixed effect and a transitory shock; the latter was assumed to follow a first order autoregressive process. The fixed effect can be thought of as being the permanent component of income. The rest is a transitory component, which may include measurement error. Their reasoning clearly originated in Friedman's work, although their emphasis was on mobility. Interestingly they also compute the ratio of the variance of the permanent component to the total variance and the number they get is similar to the one reported by Friedman. Of course such an empirical study takes an explicit stance as to what is PI; in this case it is income averaged over the observation window of the data.

The Lillard and Willis (1978) model was very influential and had implications for both consumption and for income mobility. However, an important issue has been raised since: is the permanent component of income fixed, or does it get updated by random shocks? In other words does income have a martingale component generating a stochastic trend, or are all trends deterministic? This is a key issue of course for the consumption literature because, as also implied by the PIH, any permanent shock to income would be transmitted one for one (or more) to consumption, making consumption growth at least as volatile as income growth. If it is the case that income has a unit root, then the question is whether consumption is volatile enough to be consistent with some version of the PIH. Of course it still remains the case that consumption should not respond to predictable changes in income.

A seminal paper on this issue has been written by MaCurdy (1982), where he shows that earnings from the PSID are best described by the sum of a random walk and an MA(1) component. Thus his preferred representation for income of individual i in period t can be written as

$$y_{it} = \gamma' \mathbf{x}_{it} + y_{it}^p + u_{it} \quad \text{where} \quad y_{it}^p = y_{it-1}^p + \varepsilon_{it}^p.$$

In the above ε represents a *permanent* shock while u represents a transitory shock and the \mathbf{x} represent other observed characteristics determining permanent income, such as age and education. This result has endured different estimation and testing methods and studies, and updated date sets. Abowd and Card (1989), Gottschalk and Moffitt (1993) as well as Meghir and Pistaferri (2002), find similar results and develop these ideas. This perhaps leads us to update our notion of permanent income and firmly suggests that backward looking methods of computing PI can be seriously misleading. This justifies even further Friedman's reluctance to define PI in a precise way. Combining this income process, together with the consumption model presented above, is capable of explaining a number of key observations on consumption and its distribution across individuals in a cohort.

There have been further important developments which go far beyond this short retrospective. This includes the role of precautionary savings (Deaton, 1992; Carroll, 1992; Blundell and Preston, 1998; Attanasio *et al.*, 1999) as well as the

interactions of savings and the welfare system (Hubbard *et al.*, 1995), not to mention the important and related literature on asset pricing including the seminal work by Hansen and Singleton (1983).

6. Conclusions

Most important discoveries and insights are simple, economical, have important implications for a broad range of issues and withstand the test of time. Moreover, they generate large amounts of research, verifying it and refining it. This is exactly the case with Friedman's PIH. At the end of all this, the original idea has not only survived but has formed the basis for developing a coherent analysis of consumption and savings. As such it will always be remembered as a key turning point in the development of economic science.

University College London and Institute for Fiscal Studies

References

- Abowd, J. and Card, D. (1989). 'On the covariance structure of earnings and hours changes', *Econometrica*, vol. 57(2), pp. 411–45.
- Ando, A. and Modigliani, F. (1963). 'The life cycle hypothesis of saving: aggregate implications and tests', *American Economic Review*, vol. 53(1), pp. 55–84.
- Angrist, J. and Krueger, A. (2001). 'Instrumental variables and the search for identification: from supply and demand to natural experiments', *Journal of Economic Perspectives*, vol. 15(4), pp. 69–85.
- Attanasio, O., Banks, J., Meghir, C. and Weber, G. (1999). 'Humps and bumps in lifetime consumption', *Journal of Business and Economic Statistics*, vol. 17(1), pp. 22–35.
- Attanasio, O. and Davis, S. (1996). 'Relative wage movements and the distribution of consumption', *Journal of Political Economy*, vol. 104(6), pp. 1227–62.
- Attanasio O. and Weber, G. (1993). 'Consumption growth, the interest rate and aggregation', *Review of Economic Studies*, vol. 60(3), pp. 631–49.
- Attanasio O. and Weber, G. (1995). 'Is consumption growth consistent with intertemporal optimisation? Evidence from the consumer expenditure survey', *Journal of Political Economy*, vol. 103(6), pp. 1121–57.
- Blinder, A. and Deaton, A. (1985). 'The time series consumption function revisited', *Brookings Papers on Economic Activity*, No. 2, pp. 465–511.
- Blundell, R. Browning, M. and Meghir, C. (1994). 'A microeconomic model of intertemporal substitution and consumer demand', *Review of Economic Studies*, vol. 61(1), pp. 57–80.
- Blundell, R. and Preston, I. (1998). 'Consumption inequality and income uncertainty', *Quarterly Journal of Economics*, vol. 113, pp. 603–40.
- Carroll, C. D. (1992). 'The buffer stock theory of saving: some macroeconomic evidence'. *Brookings Papers on Economic Activity*, No. 2, pp. 61–156.
- Campbell, J. (1987). 'Does saving anticipate declining labor income? An alternative test of the permanent income hypothesis', *Econometrica*, vol. 55(6), pp. 1355–65.
- Cochrane, J. H. (1991). 'A simple test of consumption insurance'. *Journal of Political Economy*, vol. 99(5), pp. 957–76.
- Deaton, A. (1992). *Understanding Consumption*, Oxford: Clarendon Press.
- Drèze, J. and Modigliani, F. (1972). 'Consumption decisions under uncertainty', *Journal of Economic Theory*, vol. 5, pp. 308–55.
- Duesenberry, J. S. (1949). *Income, Saving, and the Theory of Consumer Behavior*, Cambridge, MA: Harvard University Press.
- Eisner, R. (1958). 'The permanent income hypothesis: comment', *American Economic Review*, vol. 48(5), pp. 972–90.
- Fisher, I. (1907). *The Rate of Interest*, New York: Macmillan.
- Flavin, M. (1981). 'The adjustment of consumption to changes in expectations about future income', *Journal of Political Economy*, vol. 89(5), pp. 974–1009.
- Friedman, M. (1957). *A Theory of the Consumption Function*, Princeton: Princeton University Press.
- Friedman, M. (1992). 'Do old fallacies ever die?', *Journal of Economic Literature*, vol. 30(4), pp. 2129–32.

- Friedman, M. and Kuznets, S. (1954). *Incomes from Independent Professional Practice*, New York: National Bureau of Economic Research, 1945.
- Ghez, G. and Becker, G. (1975). 'Allocation of time and goods over the life cycle', NBER.
- Goldberger, A. S. (1972). 'Structural equation methods in the social sciences', *Econometrica*, vol. 40(6), pp. 979–1001.
- Gottschalk, P. and Moffitt, R. (1993). 'The growth of earnings instability in the US labor market', *Brooking Papers on Economic Activity*, pp. 217–72.
- Griliches, Z. (1977). 'Estimating the returns to schooling: some econometric problems', *Econometrica*, vol. 45(1), pp. 1–22.
- Hall, R. (1978). 'Stochastic implications of the life-cycle permanent income hypothesis', *Journal of Political Economy*, vol. 86(6), pp. 971–87.
- Hansen, L. P. and Singleton, K. (1983). 'Stochastic consumption, risk aversion and the temporal behaviour of asset returns', *Journal of Political Economy*, vol. 91(2), pp. 249–65.
- Heckman, J. J. H. (1974). 'Life cycle consumption and labour supply: an explanation of the relationship between income and consumption over the life cycle', *American Economic Review*, vol. 64(1), pp. 188–94.
- Hotelling, H. (1933). 'Review of "The triumph of mediocrity in business", by Horace Secrist', *Journal of the American Statistical Association*, vol. 28(184), pp. 463–5.
- Hubbard, G., Skinner, J. and Zeldes, S. (1995). 'Precautionary savings and social insurance', *Journal of Political Economy*, vol. 103(2), pp. 360–99.
- Keynes, J. M. (1935). *The General Theory of Employment, Interest, and Money*, New York: Harcourt Brace & World.
- Lillard, L. and Willis, R. (1978). 'Dynamic aspects of earnings mobility'. *Econometrica*, vol. 46(5), pp. 985–1012.
- Lucas, R. E. (1976). 'Econometric policy evaluation: a critique', in (K. Brunner and A. Meltzer, eds.), *The Phillips Curve and Labour Markets, Carnegie Rochester Series*, vol. 1, pp. 19–46.
- MaCurdy, T. (1982). 'The use of time series processes to model the error structure of earnings in a longitudinal data analysis', *Journal of Econometrics*, vol. 18(1), pp. 83–114.
- Meghir, C. and Pistaferri, L. (2002). 'Income variance dynamics and heterogeneity', *Econometrica*, vol. 72(1), pp. 1–32.
- Modigliani, F. and Brumberg, R. (1954). 'Utility analysis and the consumption function: an interpretation of cross section data', in (K. Kurihara, ed.), *Post Keynesian Economics*, pp. 388–436, New Brunswick NJ: Rutgers University Press.
- Morgan, J. N. (1951). 'The structure of aggregate personal saving', *Journal of Political Economy*, vol. 59, pp. 528–34.
- Ramsey, F. (1928). 'A mathematical theory of saving', *ECONOMIC JOURNAL*, vol. 38(152), pp. 543–50.
- Reid, M. (1952). 'Effect of income concept upon expenditure curves of farm families', in Conference on Research in Income and Wealth, *Studies in Income and Wealth*, vol. 15. New York: NBER.
- Sandmo, A. (1969). 'Capital risk, consumption and portfolio choice', *Econometrica*, vol. 37, pp. 586–99.
- Sargent, T. (1978). 'Rational expectations, econometric exogeneity and consumption', *Journal of Political Economy*, vol. 86(4), pp. 673–700.
- Wald, A. (1940). 'The fitting of straight lines if both variables are measured with error', *Annals of Mathematical Statistics*, vol. 11(3), pp. 284–300.
- Zeldes, S. (1989). 'Consumption and liquidity constraints: an empirical investigation', *Journal of Political Economy*, vol. 97(2), pp. 305–46.