Comments

'The response of consumption to income: A cross-country investigation'

by J.Y. Campbell and N.G. Mankiw

Why test the permanent income hypothesis?

John H. Cochrane

University of Chicago, Chicago, IL 60637, USA

1. Two small complaints: Spurious forecasts and the Lucas critique

The central facts driving this paper are that consumption growth is slightly forecastable, and the variables that forecast consumption growth also forecast income growth. The paper reviews a very clever interpretation of these facts in terms of 'rule of thumb' consumers that devote a fixed fraction of income to non-durable consumption.

I only have two small criticisms of the paper itself. First, the basic facts are a little tenuous. The permanent income model predicts, and the data show, a strong association between ex post consumption growth and income growth. Thus, if by fishing one were to find a variable that spuriously forecast consumption growth, that variable is practically guaranteed to also forecast income growth, and thus spuriously indicate the presence of 'rule of thumb' consumers. And the R^2 of the consumption growth forecasting regressions is low – as low as 0.01 for the U.S. – suggesting that this may be an important problem. Similarly, if one can account for the slight predictability of consumption growth by small changes in the model specification (non-quadratic or non-time separable utility, non-geometric depreciation of durables, seasonal adjustment, measurement or aggregation errors), this would also eliminate the basic facts behind the results.

Second, I have to mention the obvious methodological complaints about adopting correlations in the data as decision rules. If we can posit that some consumers follow the empirically convenient rule of thumb $c_t = y_t$, why not just estimate a vector-autoregression of consumption, income and other variables, and call the consumption equation the 'rule of thumb'?

There are two reasons we normally do not do either. First, agents in such a model ignore the utility gains of optimal behavior. Campbell and Mankiw effectively counter this criticsm by showing that their rule of thumb imposes small utility costs. Akerlof and Yellen (1985) argued for the possibility of such 'near rational' behavior, and I (1989) argued that it can represent optimal behavior with small transactions costs or other frictions.

The second reason is more important. Following the Lucas critique, consumption regressions are not invariant to regime changes. If the model is to be used for purposes other than within-regime forecasting, (and it is, otherwise it would just be a VAR with no interpretation) one must find a way to isolate structural or invariant descriptions of behavior. Showing that a given decision rule is the optimum isolates the invariant aspect of behavior. (This, and not a methodological imperative for smart agents, is the real reason for the insistence on optimizing behavior.) However, showing that a given decision rule is *near* rational does *not* show that it is invariant, since there are many non-invariant near-rational decision rules.

I have no concrete suggestions for how one should verify that a near-rational decision rule is structural, but I hope Campbell and Mankiw do, because this is probably the biggest problem blocking acceptance of near-rational rule-of-thumb methodolody. One could try to show how rules of thumb are in fact optimal when second-order costs are included, but models that include every small cost (buying the newspaper, going to the bank, etc.) are unpromising, because one can never hope to include all small costs. One hears calls for verification in micro data, but it's not clear how this will help either. Micro agents live in the same economy (regime) as the representative agent, so one can come to exactly the same (wrong) conclusions as in aggregate data. Micro data are not fundamentally more informative.

Beyond these two minor points, I think my time will be best spent reviewing the research *program* rather than Campbell and Mankiw's paper. I want to make clear at the outset that I am not critical of Campbell and Mankiw in particular. This is a fine paper in the tradition of modern empirical permanent income research. I want to quarrel with the rules of the game, not how Campbell and Mankiw have played it.

2. Why bother with the PIH?

2.1. In ISLM models

Why do we (macroeconomists) care about empirical testing of the permanent income hypothesis? Once upon a time, the answer to this question was obvious. Empirical testing of the PIH was aimed at the specification of the consumption function, which was part of a larger ISLM model. Macroeconomists essentially agreed on the form of this larger model;

disagreements centered on the form of the equations in the model. The point of the permanent income hypothesis was that the coefficient on income of the consumption function was low, and everyone knew what impact this finding would have on the behavior of the larger model.

This was part of a larger methodology, in which empirical macroeconomic research was parceled out by aggregates. By a gentleman's agreement to ignore simultaneous equations bias, one researcher could study consumption, another investment, another money demand, another labor supply, etc. Each would come up with a 'preferred specification' that one could combine to refine an ISLM model.

2.2. *Why now*

But ISLM models are dead (except, perhaps, for pure forecasting). Almost all economists now have some sort of equilibrum model in mind as the larger model. Even the most neo-Keynesian economists just advocate that the larger model should feature frictions such as menu or transaction costs, credit rationing, money, or imperfect competition, or they argue for 'sunspots' or 'bubbles' in the model's solutions. As before, we agree on the *form* of the larger model, if not its specification.

Equilibrium models are built from preferences, technology and market structure, not from consumption functions, investment schedules, labor supply equations, and so forth. One cannot study the 'compensation sector' of an equilibrium model; it does not have one. The predictions of the PIH model are not a subset of the predictions of larger equilibrium models that we can study in isolation; many equilibrium models *contradict* the predictions of the PIH. (For example, they may predict a slight forecastability of consumption changes).

Yet empirical work continues as if nothing has changed! Macroeconomists still subdivide themselves as 'consumption', 'investment', 'labor supply', etc. experts, i.e. by which of the equations of a dead model they have to work on. Top journals feature papers that test the permanent income model of consumption, and similar work on the other equations. The introductions of some recent permanent income papers even motivate the work as important to measuring the marginal propensity to consume, and studying the effects of tax cuts!

So why test the PIH? What does this empirical work tell us about larger models? I will try two answers.

2.3. *Euler equation tests*

The first component of the PIH is the Euler equation. In most PIH papers (including Campbell and Mankiw's) this prediction is that the *expected* consumption change should be zero. Since Euler equations are partial equilibrium propositions (first order conditions that lead to demand curves), they can hold no matter what the underlying production technology. Thus they *can* be studied without fully specifying the larger model, to yield information about preferences.

Yet if this is to be the reason to test the PIH, I do not see why the PIH rules apply. It is easy to study Euler equations without imposing quadratic, time- and goods-separable utility, a constant real risk-free rate, etc. These features were not introduced to satisfy an esthetic demand for generality, they help explain the data. In particular, the slight predictability of consumption growth is not a puzzle once one allows them.[1] Thus, this motivation is a good reason to work on and read the Euler equation literature, but not a rationale for the PIH.

2.4. Ex-post consumption and present value tests

The second component of the PIH is the relation between ex-post consumption changes and realizations of other variables, or the present value implications. Euler equations like (1) say *nothing* about how much Δc_t should react to news at t. Historically, the PIH evolved a separate identity from the Euler equation literature precisely to study this issue.

I can briefly summarize the ideas with a simple example. If labor income e_t follows an AR(1)

$$e_t = \rho e_{t-1} + \varepsilon_t,$$

the PIH model predicts

$$\Delta c_t = (1-\lambda) \sum_{j=1}^{\infty} \lambda^j (\mathrm{E}_t(e_{t+j}) - \mathrm{E}_{t-1}(e_{t+j})) \tag{2}$$

$$\Delta c_t = (r/(1+r-\rho)), \varepsilon_t, \quad \lambda = 1/(1+r). \tag{3}$$

In the classical case that ρ is low, (3) shows that consumption should vary less than the innovation in income. Since it varies somewhat more than this calculation suggests, the early tests found 'excess sensitivity' to income changes. It is now more fashionable to think of ρ near 1. In this case, (3) says that consumption changes should equal income changes. Since the

[1]Of course, current Euler equation tests also leave puzzles, but these are cross-sectional asset pricing puzzles such as the unconditional equity premium (which PIH tests have not yet tried to address), rather than the forecastability of consumption growth taken alone. The point is that Euler equations *can* be used to study preferences, not that they have yet come to a definitive answer.

standard deviation of consumption growth is about 1/2 that of income changes, consumption is 'excessively smooth'.

I have two reservations about this motivation for the PIH. First, Hansen, Roberds and Sargent (1990) (henceforth HRS) show that the present value implications are not testable, because they are not robust to the possibility that agents have more information than we do. Their argument runs roughly as follows. We start by stating what the PIH restrictions are. Given that income follows

$$\Delta e_t = \rho(L) w_t, \tag{4}$$

(w_t may be multidimensional), the PIH (2) can be written

$$\Delta c_t = \rho(\lambda) w_t, \tag{5}$$

[see Sargent (1987)]. We can rewrite (4) as a forecast of income using current and lagged consumption changes and other variables:

$$\Delta e_t = \alpha(L)\Delta c_t + \beta(L)\varepsilon_t. \tag{6}$$

Then, the PIH (5) implies

$$\Delta c_t = \alpha(\lambda)\Delta c_t + \beta(\lambda)\varepsilon_t. \tag{7}$$

Hence,

$$\alpha(\lambda) = 1, \quad \beta(\lambda) = 0.$$

The first restriction says that if consumption changes by \$1, this should signal a \$1 change in the present value of future income; the second says that if ε_t changes but consumption did not change, the present value of subsequent changes in future income should be zero.

These look like testable restrictions, but HRS show they are not. If $\beta(\lambda) \neq 0$, HRS show that one can always construct a larger information set that consumers can be imagined to observe (but we do not), such that $\beta(\lambda) = 0$ again with respect to this larger information set. For example, suppose agents see

$$v_t = D(L^{-1})\varepsilon_t = ((L^{-1} - \lambda)/(1 - \lambda L^{-1}))\varepsilon_t.$$

One can verify that $D(z)D(z^{-1}) = 1$, so v_t is serially uncorrelated, and $D(L^{-1}) = D(L)^{-1}$, so $\varepsilon_t = D(L)v_t$. By construction, $D(\lambda) = 0$. Proceeding as above, the income eq. (6) is

$$\Delta e_t = \alpha(L)\Delta c_t + \beta(L)D(L)v_t,$$

so the PIH (7) predicts

$$\beta(\lambda)D(\lambda) = 0,$$

which is true by construction. Therefore, one might as well leave other variables out of the income forecasting regression. In particular, all the arguments about how lagged income forecasts future income are irrelevant.

HRS also show that $\alpha(\lambda)=1$ is not testable. Unless non-durable consumption goods are the *only* consumption goods, one can construct a sequence of economies, each with $\alpha(\lambda)=1$, that approximate a given economy with $\alpha(\lambda)\neq 1$ arbitrarily well. Furthermore, HRS show that the restriction $\alpha(\lambda)=1$ does not even hold when one includes unanticipated capital gains, as almost all empirical studies do.

My second reservation is that the present value implications are unlikely to survive when the PIH model is used in a larger model. There are two ways to see this point. First, note that the budget constraint implications require a specification of *all* the assets available to the consumer.[2] The Euler equation implications hold for each asset separately, independently of how many assets are really out there. Thus, the budget constraint implications do not survive when consumers are placed in an environment with a richer set of assets than a risk-free rate.

Second, we can interpret the PIH model as a completely specified equilibrium model, rather than a part of some larger model. To see this, note that the PIH model can be written as

$$\max \mathrm{E} \sum_{t=0}^{\infty} \lambda^t(-(c_t-c^*)^2) \quad \text{s.t.} \begin{cases} y_t = mpk\, k_t + e_t, \\ y_t = c_t + i_t, \\ k_{t+1} = (1-\delta)k_t + i_t, \end{cases}$$

$\lambda = 1/R, \qquad R = 1 + mpk - \delta.$

Plus a bound or transversality condition on k to rule out Ponzi schemes. The planning problem of a very simple equilibrium model might be

[2]This is technically equivalent to specifying the production technology in complete markets, so this interpretation and the one that follows are not really all that different.

$$\max \mathrm{E} \sum_{t=0}^{\infty} \lambda^t u(c_t\, l_t) \quad \text{s.t.} \begin{cases} y_t = f(k_t, l_t, \text{shock}_t), \\ y_t \quad = c_t + i_t, \\ k_{t+1} = (1-\delta)k_t + i_t. \end{cases}$$

Plus a similar bound.

The PIH model is just a special case of the equilibrium model, not its 'consumption sector'. It is not 'placed in' the larger model, it is 'generalized to' the larger model, and its implications are unlikely to survive this generalization.

One might take this interpretation literally, and read tests of the PIH as tests of a specific equilibrium model. [This seems to be the approach advocated by Eichenbaum, Christiano and Marshall (1990).] But why? The state of the art in stochastic equilibrium models has advanced far beyond quadratic utility and linear technology. Again, this generalization was not undertaken to satisfy an esthetic demand for complexity, but in order to provide a better match to the data; and not only in the attempt to explain one set of moments (the second moments of consumption and labor income) but moments of *all* variables. More complex equilibrium models have been constructed *precisely because* the PIH model does describe the data.

One might defend the PIH model as an *approximation* to the larger model, but then why does a rejection suggest 'liquidity constraints' or 'rule of thumb behavior' rather than (now technically simple) relaxation of the approximations? If the PIH model is a simple-to-solve, pedagogically useful approximations to the true model, why are we not satisfied that it is 'approximately' correct, as it certainly is?

3. Summary and conclusions

I posed the question, 'Why bother sutrying the PIH?' or, more precisely, 'How does the empirical literature that tests the PIH help us to specify or refine equilibrium models?' (What else could its purpose be?) I conclude that it does not.[3] Its Euler equation implications – that consumption changes should not be forecastable – do survive in larger models, but they are easily generalized beyond quadratic utility and constant real interest rates. Its present value implications – the reponse of consumption to income or other shocks – are not testable, and do not survive in larger models. The same points hold for the stylized equilibrium models that are the heirs of empirical work on the investment, labor supply, money demand, etc. equations.

[3]It is not very good at improving the specification of ISLM models either. The modern PIH is essentially intertemporal, and does not yield a 'consumption function', or any other object that one could easily insert into an ISLM model. Neither do the λ model, or other competitors, such as models that feature explicit 'liquidity constraints'.

So what should we do instead? Work on equilibrium models can be split up into independent pieces, but by *agents* rather than by *aggregates*: One researcher can study *households*' consumption and labor supply (simultaneously) and infer their preferences. Another can study *firms*' output supply, investment and labor demand and infer their technology. Euler equation methods – essentially measures of marginal rates of substitution – are particularly convenient to these studies, as they do not require a complete specification of each agent's budget constraint, or explicit solution for his decision rule. These investigations could go on independently of each other, and contribute information that survives when placed in an equilibrium model, even one with frictions, and is therefore useful to the refinement of such models.

If we had never seen ISLM models, we probably would have divided empirical work this way rather than continuing and trying to adapt work on ISLM equations. In fact, equilibrium model-builders, tired of having to re-estimate taste and technology in each paper to match equilibrium dynamics with data, are already adopting the preferences from the latest Euler equation tests. Thus, I suspect that this pattern of work will emerge spontaneously, though later than one might have predicted.

References

Akerlof, George A. and Janet L. Yellen, 1985, A near rational model of the business cycle with wage and price inertia, The Quarterly Journal of Economics 100, 824–838.

Cochrane, John H., 1989, The sensitivity of tests of the intertemporal allocation of consumption to near rational alternatives, American Economic Review 79 (June) 319–337.

Eichenbaum, Martin S., Lawrence Christiano and David Marshall, 1990, The permanent income hypothesis revisited, Econometrica, forthcoming.

Hansen, Lars Peter, William Roberds and Thomas J. Sargent, 1990, Time series implications of present value budget balance and of martingale models of consumption and taxes, Manuscript.

Sargent, 1987, Macroeconomic theory, 2nd ed. (Academic Press, New York).

Comments

'The response of consumption to income: A cross-country investigation'

by John Y. Campbell and N. Gregory Mankiw

David F. Hendry

Nuffield College, Oxford, OX1 1NF, UK

The paper by John Campbell and Gregory Mankiw is an interesting