Stochastic Implications of the PIH: A Critique

Thomas von Ungern-Sternberg*

April 20, 2018

Abstract

This note critically studies the Hall (1978) paper testing the Permanent Income Hypothesis (PIH). We note that: (i) the econometric methodology used by Hall is unconventional at best; (ii) if the classical version of the PIH holds, observed consumption will NOT follow a random walk; (iii) depending on the stochastic properties of observed consumption and income one may obtain the results that either the test proposed by Hall systematically rejects the PIH even though it is in fact correct or the test systematically accepts the hypothesis even though it is incorrect.

Keywords: Permanent Income Hypothesis, Measurement Error

JEL Codes: ?

*Faculté des HEC, UNIL

1 Introduction

The purpose of this note is to critically examine the well-known (4300 Google-, 1200 SSCI-citations) paper by Hall (1978), which purports to test the stochastic implications of the Permanent Income Hypothesis (PIH). The test can be briefly summarized as follows: Suppose you have a regression of C_t on C_{t-1}. If the PIH holds, then adding further lags of C_t and/or lagged values of Y_t (Y_{t-1}, Y_{t-1}, etc) to this regression should not increase its explanatory power. The coefficients of these terms should not be significantly different from zero.

Three points about this test will be highlighted. First, the econometric methodology underlying this test is unconventional/doubtful. Typical statistical testing is based on the idea that one accepts the hypothesis that a coefficient (or a group of coefficients) is different from zero, only if there is a less than $x\%$ chance of accepting the hypothesis even though it is wrong. The Hall methodology is to hypothesize that a coefficient should be equal to zero, and then accept the hypothesis if it is not significantly different from zero. An analogy would be to hypothesize that a coefficient should be equal to say 50 and then accept the hypothesis if the estimate is not significantly different from 50 (i.e. not even ask the question whether the coefficient is in fact significantly different from zero). Alternatively formulated: with Halls methodology, there is no way to determine how high is the probability that one ends up accepting a false hypothesis.

Second, we argue that the basic methodology of Hall is fundamentally flawed. Even if the PIH is in fact correct, it is perfectly possible that past levels of income do in fact significantly improve the explanatory power of the regression. The property on which Hall bases his test is valid only under extremely restrictive (not to say unrealistic) implicit assumptions.

Third, we will argue that it is very easy to generate results where the stochastic property Hall tests for does in fact hold...even though the data are in fact generated by a simple distributed lag distribution model.

In summary, the property Hall bases his test on is neither a necessary nor a sufficient condition for the PIH to hold.

2 The basic model

To understand these points, let us start with a strongly simplified version of the PIH.

$$C_t^P = k \cdot Y_t^P \tag{1}$$

$$Y_t^P = (1+g) \cdot Y_{t-1}^P + \epsilon_t \tag{2}$$

where the period is denoted by index t, C^P is permanent consumption, Y^P is permanent income, k is the marginal propensity to consume, g is a growth rate and ϵ_t is a serially uncorrelated random shock. Since $Y_{t-1}^P = \frac{1}{k} \cdot C_{t-1}^P$, one can easily rewrite (1) in the form

$$C_t^P = (1+g) \cdot C_{t-1}^P + k \cdot \epsilon_t \tag{3}$$

This is the basis of Hall's approach. According to the PIH, consumption in time t is just equal to (a multiple of) consumption in time $t-1$ plus a change due to new information about the evolution of permanent income. Permanent consumption should be a random walk with drift.

To understand the limits of this approach it is useful to explicitly distinguish between the theoretical concepts C^P and Y^P and the variables one actually observes C and Y. The differences between the two may be due to measurement errors and/or transitory components. We will adopt the latter formulation, though they are (in our model) equivalent.

According to the PIH as formulated by Friedman (1957), the consumption expenditures one actually observes are the sum of the permanent plus the transitory component $C_t = C_t^P + C_t^T$. In its simplest form, it is assumed that transitory consumption C_t^T is white noise. Similarly, it is assumed that $Y_t = Y_t^P + Y_t^T$, where Y_t^T is also white noise. When one regresses observed C_t on observed C_{t-1} one regresses $(C_t^P + C_t^T) = \gamma \cdot (C_{t-1}^P + C_{t-1}^T) + \epsilon_t$

In this formulation it is easy to see that the estimate of gamma one obtains will NOT in general be equal to (1+g) but lower. There will be attenuation bias, due to the appearance of transitory consumption on the RHS. The important thing to note is that one can obtain an unbiased estimate of $1+g$ only if one makes the unrealistic assumption that the observed Cs are equal to C^P. (Hall does in fact make this assumption...implicitly as several authors have noted. No justification or discussion is provided for this extreme assumption).

Note that even if one were willing to assume that observed total consumer expenditure is in fact equal to permanent consumption (i.e. transitory consumption is identically equal to zero) this would still NOT justify Halls methodology. The data he uses are not total consumption expenditure but expenditures on services and non-durables (i.e. he excludes durable expenditure). Blinder and Deaton (1985) have shown that the relative price of durables to nondurables has a significant effect on expenditure shares. The conclusion: even if total consumer expenditure is in fact a random walk with drift, the same should not generally be true for nondurable consumer expenditure (there is no reason for relative price changes to be uncorrelated over time).

3 ?? Name ??

Let us now study what would be the effect of adding the term Y_{t-1} to a regression of C_t on C_{t-1}. Essentially, what one would be doing is to regress C_t on two noisy signals of Y^P_{t-1}. In case the PIH does hold, the relative size of the coefficients should depend on the degree of noisiness of the two signals. If transitory consumption is very low, one should expect the coefficient of C_{t-1} to be close to one, and the coefficient of lagged income to be close to zero. If transitory income is small, one would expect the coefficient of Y_{t-1} to be close to one, and the coefficient of lagged consumption to be close to zero. Depending on the relative magnitude of the transitory components, all kinds of intermediate results can of course be generated.

Halls basic hypothesis that only past consumption should matter is thus incorrect. Or to be more precise, it is correct only under the very extreme assumption that consumption has no transitory component. For real world data, the fact that Y_{t-1} has a significant coefficient cannot be used as an indication that the PIH should be rejected.

Illustrate this with simulation results??

The presence of transitory components in observed income may also be the reason why one might accept the PIH even though the data are in fact generated by a standard distributed lag model. To illustrate this effect, we have generated the following very simple model, where by construction:

$$C_t^P = 0.3 \cdot Y_t^P + 0.3 \cdot Y_{t-1}^P + 0.3 \cdot Y_{t-2}^P + \eta_t \tag{4}$$

(permanent consumption does not instantly adapt to changes in permanent income, but is spread over three periods). Y_t^P is a simple random walk (with drift):

$$Y_t^P = 1.005 \cdot Y_{t-1}^P + \mu_t \tag{5}$$

When the transitory component of Y is low relatively to that of C, one obtains:

$$C_t = \mathit{constant?} + \underset{(216.30)}{0.29} \cdot Y_t + \underset{(216.30)}{0.30} \cdot Y_{t-1} + \underset{(216.30)}{0.30} \cdot Y_{t-2} \tag{6}$$

$$C_t = \underset{(216.30)}{3} + \underset{(216.30)}{0.96} \cdot C_{t-1} \tag{7}$$

$$C_t = \underset{(216.30)}{0.5} + \underset{(216.30)}{0.14} \cdot C_{t-1} + \underset{(216.30)}{0.45} \cdot Y_{t-1} + \underset{(216.30)}{0.32} \cdot Y_{t-2} \tag{8}$$

As was to be expected, the lagged values of Y do most of the explaining. When the transitory component of Y is high, the situation changes substantially:

$$C_t = \underset{(216.30)}{6} + \underset{(216.30)}{0.26} \cdot Y_t + \underset{(216.30)}{0.26} \cdot Y_{t-1} + \underset{(216.30)}{0.26} \cdot Y_{t-2} \tag{9}$$

$$C_t = \underset{(216.30)}{3.6} + \underset{(216.30)}{0.96} \cdot C_{t-1} \tag{10}$$

$$C_t = \underset{(216.30)}{3.3} + \underset{(216.30)}{0.85} \cdot C_{t-1} + \underset{(216.30)}{0.04} \cdot Y_{t-1} + \underset{(216.30)}{0.05} \cdot Y_{t-2} \tag{11}$$

Hall would interpret the results in equation 11 as giving support to the PIH. In reality, the numbers were generated by a standard distributed lag process. The driving force underlying the results is the transitory component in total income. It should also be emphasized that the results are not simply due to the standard attenuation bias (measurement errors in the RHS variable lead to lower coefficients). As the first equation shows the standard attenuation bias is only of the order of around 10% (from 0.3 to 0.26). The problem is greatly magnified by the misspecification, i.e. the introduction of a lagged values of C_t on the RHS. As these lagged value of C_t is not influenced by the transitory components of Y, it ends up doing most of the explaining. In such a misspecified equation, one can easily fall prey to the illusion that the past values of income play no significant role.

4 Cyclical behaviour

Hall writes: This implication would be refuted if consumption had a definite cyclical pattern described by a difference equation of second or higher order. Intelligent consumers ought to be able to offset any such cyclical pattern and restore the noncyclical optimal behavior of consumption predicted by the hypothesis.

In reality consumer expenditure does typically have a strong cyclical pattern. In the western world expenditure is typically highest in the 4th quarter (Xmas effect) and lowest in the 1st quarter (hang-over effect?). If Halls restrictive interpretation of the PIH were in fact correct, the theory would never had any success. There is no theoretical reason to say that all expenditures on Xmas gifts must be offset in the same quarter.

What Hall does in fact test for is actually quite surprising. He uses cyclically adjusted data on consumer expenditure and then tests whether these data have a cyclical pattern??? If they did, the only conclusion one could draw is that the US statisticians are not be doing a particularly good job!!!

5 How robust are Hall's results ?

We have run some regressions using Halls specification with US data for the longer time period 1948- 2015.

Using annual data we obtain:

$$C_t = \underset{(216.30)}{690.25} + \underset{(12.58)}{36.19} \cdot t + \underset{(0.08)}{1.07} \cdot C_{t-1} - \underset{(0.06)}{0.13} \cdot Y_{t-1}$$

The hypothesis that the coefficient of Y_{t-1} is equal to zero should be rejected. With quarterly (adjusted) data the results are:

$$C_t = \underset{(39.28)}{143.24} + \underset{(0.56)}{1.77} \cdot t + \underset{(0.01)}{1.02} \cdot C_{t-1} - \underset{(0.01)}{0.03} \cdot Y_{t-1}$$

Once again the hypothesis that the coefficient of Y_{t-1} is equal to zero should be rejected. It should be noted that these results are not really interesting. As explained above both the statistics and the economic theory underlying the tests are faulty. The one regularity we have observed is that the coefficient of the Y_{t-1} term is typically small and negative. Furthermore (as the regressions above show) the coefficient is typically smaller for quarterly data than for annual data. Hall interprets the small coefficient of the Y_{t-1} term as follows: There is a statistically marginal and numerically small relation between consumption and very recent levels of disposable income (p 984). He has no explanation why the coefficient of the Y_{t-1} term should be negative attributing it to sampling variation.

There is an alternative very simple explanation. Suppose that:

$$Ct = a \cdot Y_t + (a-d) \cdot Y_{t-1} + (a-2d) \cdot Y_{t-2} + (a-3d) \cdot Y_{t-3} + ...$$

till the sum of the marginal propensities to consume is just equal to 1. We then have :

$$C_t - C_{t-1} = a \cdot Y_t - d \cdot Y_{t-1} - d \cdot Y_{t-2} - ...$$

Excluding the Y_t term (as Hall does) we have:

$$C_t = C_{t-1} - d \cdot Y_{t-1} - d \cdot Y_{t-2} - ...$$

This offers an alternative (more plausible) explanation of the estimates one obtains. First the small coefficient of the Y_{t-1} term does not indicate that recent levels of disposable income have a small influence. To the contrary: it indicates that past levels of disposable income have roughly the same effect as current levels. The reason the coefficient of Y_{t-1} is negative is not due to sampling variation. It is due to the fact that the impact effect of a change in income is larger than the lagged effect.

6 Conclusion

If one accepts that observed consumption does in fact contain a transitory component (and there is neither empirical nor theoretical evidence that it does not) than the test of the PIH Hall proposes can give neither a necessary nor a sufficient condition for the PIH to be valid. Furthermore he uses an econometric methodology that should be avoided whenever possible. A modern statistical test should give some indication how high is the probability of accepting a false hypothesis.