

Income Asymmetries and the Permanent Income Hypothesis^{*}

Juan I Urquiza[†]
University of Houston

November 2010

Abstract

Within the context of the Permanent Income Hypothesis (PIH), the predictions for consumption depend crucially upon the process for income. In this paper, we consider an unobserved components model that allows for both asymmetric transitory movements and correlation between permanent and transitory innovations. Using aggregate U.S. data, we show that this model fits labor income data significantly better than common alternatives. However, we find that consumption is excessively smooth relative to the predictions of our model. To reconcile these predictions with the data, we explore the possibility of imperfect information. A delayed information version of the model fits the data better but consumption is excessively sensitive compared to the predictions of this model. We are able to match the data when we consider an economy in which 60 – 65% of consumers behave according to the PIH with full information and the remaining consumers have delayed information.

KEYWORDS: Unobserved Components Models; Markov-Switching; Consumption Dynamics; Excess Smoothness; Excess Sensitivity; Imperfect Information.

JEL CLASIFICATIONS: C22, C51, E21, E32.

^{*} I owe special thanks to Tara Sinclair for sharing her code. I am also grateful to James Morley, Chris Murray, David Papell, Bent Sørensen, and Dietrich Vollrath for useful comments.

[†] University of Houston, Department of Economics, 204 McElhinney Hall, Houston, TX 77204-5019. E-mail: jiurquiza@uh.edu.

I. Introduction

In forward-looking models like Hall’s (1978) version of the Permanent Income Hypothesis (PIH), the predictions for consumption depend crucially upon the process describing income. For example, Hansen and Sargent (1981) show that when income is represented by an autoregressive moving average (ARMA) process, the model provides closed-form solutions for the predicted change in consumption due to income innovations. Conditional on this representation, consumption growth should be more volatile than income growth if income follows a highly persistent process – as suggested by quarterly data. Yet, studies using aggregate data have consistently found that consumption growth is much smoother than income growth. In the literature, this phenomenon is referred to as “excess smoothness.”

In this paper, we derive and test the predictions of the PIH when income is represented by an unobserved components (UC) model. Typically, it is assumed that income can be decomposed into the sum of a permanent random walk component and a stationary transitory component. However, in the light of recent developments in the business cycle literature, we extend the basic model to allow for asymmetric transitory movements and correlation between permanent and transitory innovations as in Sinclair (2010).¹ By introducing these modifications, we are able to address a key concern within the consumption literature: i.e., the relative importance of permanent shocks. Intuitively, if the transitory component of income is asymmetric in the sense that the mean growth rate of income during recessions differs significantly from that during expansions, then a linear symmetric model may over-emphasize permanent movements due to the predominance of expansions in the data.

¹ Sinclair (2010) shows that in the context of U.S. quarterly real GDP, ignoring the correlation between innovations underestimates the role of permanent movements, while ignoring asymmetry in the transitory component underestimates the role of temporary movements.

Correct identification of permanent and transitory income shocks is crucial in order to assess the validity of the PIH.² Figure 1 highlights the relevance of capturing the effect of recessions in order to understand consumption patterns. In particular, we observe a strong correlation between the growth rate of consumption and the National Bureau of Economic Research (NBER) recession dates. Importantly, if recessions are predominantly characterized by transitory negative shocks as emphasized in the business cycle literature, then the standard version of the PIH predicts that the change in consumption should be small. However, it is evident that consumption systematically drops in the event of a recession. While buffer stock savings models can provide an explanation for this behavior, we argue that the PIH is also capable of explaining these episodes when coupled with a broader specification of the process describing income.³

[Insert Figure 1]

Our results show that even when controlling for correlation and asymmetric transitory movements, the PIH fails to fit the data: the marginal propensity to consume (MPC) out of permanent shocks is significantly different from one, and consumption growth is more volatile than income growth. Therefore, we extend our model to allow for imperfect information as considered by Goodfriend (1992) and Pischke (1995). In particular, we assume that agents know the structure of the economy but cannot distinguish between permanent and transitory shocks at time t ; this information becomes available with one period lag. In this framework, the model's predictions are in line with the data both in terms of MPCs and smoothness but consumption

² See Flavin (1981) and Hall and Mishkin (1982).

³ See Carroll (1992).

appears to be too sensitive.⁴ It is only when we allow for the economy to be populated by agents with different information sets that we are able to match the data.

This paper relates to Quah (1990), who shows that if the econometrician observes income news different from the news that individuals observe, then she might reject the PIH even when individuals behave according to it. He estimates different UC models for various reduced-form ARMA models describing aggregate income, and provides a solution to the excess smoothness puzzle by showing that there always exists an UC model that makes the PIH consistent with the data. However, he does not address the issue of whether those decompositions are reasonable, nor does he attempt to explicitly identify the structural disturbances that affect labor income. That is precisely what we try to do in this paper.

Our paper also relates to Hryshko (2008). He uses simulation experiments to show that allowing for correlation between permanent and transitory shocks can be very important for interpreting lifecycle consumption. In particular, he finds that household level data can be better fitted when income shocks are negatively correlated. However, despite the fact the sum of innovations is smoother compared to that in income models that feature uncorrelated or positively correlated shocks, his buffer stock model of savings is unable to explain excess smoothness unless augmented to allow for partial risk sharing as in Attanasio and Pavoni (2007).

Finally, this paper relates to Goodfriend (1992) and Pischke (1995) who consider the role of imperfect information. In particular, they assume that individuals observe their own income but have little or no information about aggregate income. As a consequence, individuals cannot distinguish between aggregate (permanent) and idiosyncratic (transitory) shocks and therefore

⁴ The PIH predicts that consumption is martingale and that consumption growth should be orthogonal to predictable income changes. However, regressions using aggregate data typically find a small but significant correlation. This phenomenon is known as “excess sensitivity.”

fail to appropriately adjust their consumption. However, if information on aggregate income becomes available in subsequent periods, consumption will be revised and further adjustments will be required in order to restore the optimal plan. Thus, observed consumption will not only appear to be too smooth but also excessively sensitive to lagged income changes.

The rest of the paper is organized as follows. In section II, we present the process for income and we derive the implications of the PIH. Section III examines the performance of our model for income relative to common alternatives and shows that our representation fits the data significantly better. In section IV, we use the estimated income process to identify permanent and transitory shocks, and then we test for the PIH. The results show that our model is still incapable of replicating the smoothness of consumption and that marginal propensities to consume are far from those outlined by the theory. Section V considers the possibility of delayed information and while we are able to match the MPCs and the observed smoothness, consumption appears to be excessively sensitive. Only when we allow for the economy to be populated with consumers that differ in their respective informational assumptions, we are able to match the data. Section VI concludes.

II. The Model

Hall (1978) shows that under strict assumptions on preferences and technology (i.e., agents are forward-looking, have rational expectations and quadratic utility, and there is free borrowing and lending at a constant interest rate) the change in consumption, $\Delta \tilde{c}_t$, is given by the following expression:

$$\Delta \tilde{c}_t = \left(\frac{r}{1+r} \right) \sum_{k=0}^{\infty} (1+r)^{-k} (E_t - E_{t-1}) y_{t+k}. \quad (1)$$

Equation (1) implies that given a particular process for income, the response of consumption to income innovations can be easily calculated as the annuity value of the revisions in expected income. Therefore, we now turn to the proposed model for income.

A. The Income Process

Following Sinclair (2010), we assume that labor income can be decomposed into the sum of two unobserved components:

$$y_t = \tau_t + c_t, \quad (2)$$

where τ_t represents the permanent component and c_t represents the transitory component.

To be consistent with the literature, we assume that the permanent component follows a random walk with drift:

$$\tau_t = \mu + \tau_{t-1} + v_t, \quad (3)$$

where v_t is a permanent shock.

The transitory component is modeled as an AR(2) process, where the innovation is assumed to be a mixture of a symmetric transitory shock, u_t , and a discrete, asymmetric transitory shock, πS_t , of the form:

$$c_t = \phi_1 c_{t-1} + \phi_2 c_{t-2} + u_t + \pi S_t. \quad (4)$$

The innovations v_t and u_t are assumed to be jointly normally distributed random variables with zero mean and covariance matrix given by the following expression:

$$\Sigma = \begin{bmatrix} \sigma_v^2 & \sigma_{vu} \\ \sigma_{vu} & \sigma_u^2 \end{bmatrix} \quad (5)$$

The way we model the transitory component is the key to this paper. Contrary to most of the existing literature, we do not allow for MA terms.⁵ We do this, however, so as to satisfy Morley, Nelson, and Zivot's (2003) order condition for identification. They show that when the transitory component follows a stationary and invertible ARMA (p, q) process, there will be at least as many nonzero autocovariance relations as parameters to estimate if $p \geq q + 2$. Therefore, allowing for a MA process would require a higher order AR process in order to identify the correlation between innovations.

Our main contribution, however, is the introduction of a form of asymmetry in the transitory component. As in Hamilton (1989) and Kim and Nelson (1999), we assume that the unobserved variable, S_t , evolves according to a first-order Markov-switching process:

$$\Pr[S_t = 0 | S_{t-1} = 0] = p_{00}, \quad (6)$$

⁵ See MaCurdy (1982), Hall and Mishkin (1982), and Pischke (1995).

$$\Pr[S_t = 1 | S_{t-1} = 1] = p_{11}. \quad (7)$$

This variable can be thought of as the state of the economy, and it is endogenously determined by the model. The intuition is very straightforward: during normal times, $S_t = 0$ and labor income is near the trend; during recessions, $S_t = 1$ and the economy is hit by a negative transitory shock – or “pluck” – that pulls labor income away from the trend.⁶ In this case, identification requires that we restrict the sign of the π so we follow Sinclair (2010) and impose a non-positive constraint.

It is worth mentioning that this model nests both the symmetric and uncorrelated models as special cases (i.e., $\pi = 0$ and/or $\sigma_{vu} = 0$). Section III estimates these models and compares their relative performances.

B. Predictions of the PIH

Conditional of this process for income, and on the assumption that agents can distinguish the separate components of income given in equation (2) while observing the overall state of the economy, S_t , it can be shown that the PIH implies that:

$$\Delta \tilde{c}_t = v_t + \left(\frac{r}{1+r} \right) \cdot \psi \left(\frac{1}{1+r} \right) u_t + \gamma_{S_t}, \quad (8)$$

where $\psi(L) = (1 - \phi_1 L - \phi_2 L^2)^{-1}$ is a polynomial in the lag operator.⁷ In other words, consumption follows a martingale difference sequence with switching drift, γ_{S_t} .

Overall, equation (8) is no different from that in the literature: consumption responds one to one to permanent shocks and it is nearly insensitive to transitory shocks. However, conditional

⁶ See Friedman (1969, 1993).

⁷ A detailed derivation of equation (8) is available in appendix A (to be added).

on our model for income, the prediction of the PIH explicitly allows for a switching drift term. Given the current and previous state of the economy (i.e., S_t and S_{t-1}), it can be shown that the drift term adopts the following form:

$$\gamma_{00} = \left(\frac{r}{1+r}\right) \cdot \psi\left(\frac{1}{1+r}\right) \cdot \left[\frac{\pi \cdot (1+r)}{2+r-p_{00}-p_{11}}\right] \cdot (p_{00}-1), \quad (9.1)$$

$$\gamma_{01} = \left(\frac{r}{1+r}\right) \cdot \psi\left(\frac{1}{1+r}\right) \cdot \left[\frac{\pi \cdot (1+r)}{2+r-p_{00}-p_{11}}\right] \cdot (-p_{11}), \quad (9.2)$$

$$\gamma_{10} = \left(\frac{r}{1+r}\right) \cdot \psi\left(\frac{1}{1+r}\right) \cdot \left[\frac{\pi \cdot (1+r)}{2+r-p_{00}-p_{11}}\right] \cdot (p_{00}), \quad (9.3)$$

$$\gamma_{11} = \left(\frac{r}{1+r}\right) \cdot \psi\left(\frac{1}{1+r}\right) \cdot \left[\frac{\pi \cdot (1+r)}{2+r-p_{00}-p_{11}}\right] \cdot (1-p_{11}). \quad (9.4)$$

Notice that $(2+r-p_{00}-p_{11}) > 0$ by definition, and that $\pi \leq 0$ by assumption, so the sign of γ_{S_t} ultimately depends on the last term of equation (9) as $\left(\frac{r}{1+r}\right) \cdot \psi\left(\frac{1}{1+r}\right) \cdot \left[\frac{\pi \cdot (1+r)}{2+r-p_{00}-p_{11}}\right]$ is non-positive. Therefore, we can conclude that γ_{00} and γ_{01} are positive, while γ_{10} and γ_{11} are negative since $0 \leq p_{ii} \leq 1$ by definition.

Some interesting business cycle implications emerge as we consider the estimates for p_{00} and p_{11} . In particular, it appears that $\gamma_{00} < \gamma_{01}$ which suggests that consumption grows faster than average after a recession. This is consistent with the idea of a third phase of the business cycle introduced by Beaudry and Koop (1993) and by Sichel (1994). Basically, they find evidence supporting the existence of three instead of two phases of the business cycle: a normal phase, a recessionary phase, and a high-growth recovery phase during which output reverts to its previous peak so that the effect of recessions is transitory.

We find that there are at least two things that make equation (8) appealing. First, the fact that $\gamma_{00} > 0$ could potentially account for the observed upward trend in aggregate consumption. Second, and perhaps more important, the fact that the drift switches to γ_{10} in the event of a recession can help explain why consumption systematically drops, even if these are predominantly caused by transitory shocks.

III. Data and Maximum Likelihood Estimation

A. Data

We use U.S. seasonally adjusted quarterly data from 1952:1 to 2009:4; a total of 232 observations. Labor income data comes from Ludvigson's website, and it is compiled from the National Income and Product Accounts (NIPA) tables.⁸ For consumption, we consider three alternative measures: consumption expenditure in (i) nondurables and services, (ii) nondurables, and (iii) services. We do this so as to address two points made in the literature: first, Carroll and Sommer (2003) emphasize measurement error problems in quarterly services data and, therefore suggest excluding them from the analysis altogether; second, Ludvigson and Michaelides (2001) note that most of the excess smoothness comes from services expenditure.⁹ For comparison, we replicate Ludvigson and Michaelides' (2001) Table 1.¹⁰

[Insert Table 1]

⁸ <http://www.econ.nyu.edu/user/ludvigsons/>

⁹ For example, the imputed rent on housing is constructed using interpolated data from an annual survey of house prices.

¹⁰ Throughout this paper, we adopt their definitions for the smoothness ratio and the sensitivity coefficient. Basically, the former refers to the ratio of the standard deviation of consumption growth to that of labor income growth, while the latter is given by the ordinary least squares coefficient from a regression of consumption growth on lagged income growth.

Table 1 shows that the standard deviation of consumption growth is about half the standard deviation of labor income growth when the measure of consumption includes services expenditure. However, this ratio rises to 0.81 for expenditure on nondurables, which confirms that much of the observed smoothness comes from services. Moreover, consumption growth is positively correlated with lagged income growth, with an estimated sensitivity coefficient that ranges from 0.09 to 0.22. Contrary to the effect on the smoothness ratio, including services appears to strengthen the correlation between consumption growth and lagged income growth.

We transform labor income and consumption series into the natural logarithm of real per capita variables times 100. The data is expressed in billions of chain-weighted 2005 dollars.

B. *Maximum Likelihood Estimation*

To estimate the model presented in section II, it is cast into the following state-space form:

$$\text{State Equation:} \quad \begin{bmatrix} \tau_t \\ c_t \\ c_{t-1} \end{bmatrix} = \begin{bmatrix} \mu \\ \pi S_t \\ 0 \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & \phi_1 & \phi_2 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \tau_{t-1} \\ c_{t-1} \\ c_{t-2} \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} v_t \\ u_t \end{bmatrix} \quad (10)$$

$$\text{Observation Equation:} \quad y_t = \begin{bmatrix} 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} \tau_t \\ c_t \\ c_{t-1} \end{bmatrix} \quad (11)$$

The parameters and the unobserved components are estimated using Sinclair's (2010) code.¹¹ In essence, she applies Kim's (1994) method of combining Hamilton's (1989) algorithm and a nonlinear discrete version of the Kalman filter as an approximation to maximum likelihood estimation.

¹¹ <http://home.gwu.edu/~tsinc/>

Table 2 reports the maximum likelihood estimates of the parameters and their standard errors. Notice that models 2 and 3 are restricted versions of model 1. Some notable differences between the models are worth mentioning. First, we observe that when comparing models 1 and 2 with 3, the sum of AR coefficients for the transitory component is relatively smaller, which suggests that the persistence of the transitory component is reduced once we allow for asymmetry. Second, the asymmetric shock parameter, π , is highly significant in both of these models, and including it appears to represent an improvement over the symmetric model. Finally, it seems that ignoring either asymmetry or the correlation between innovations, ρ_{uv} , can seriously affect the estimates of the standard deviation of both permanent and transitory shocks.

[Insert Table 2]

It is important to note that the estimates of the correlation reported in Table 2 are negative and significant for both UC-UR models. This finding is consistent with Watson (1986), Stock and Watson (1988), and Morley, Nelson, and Zivot (2003). Intuitively, positive permanent shocks shift the long run path of income so that short term fluctuations largely reflect adjustments towards the shifting trend. This implies a negative contemporaneous correlation until actual income catches up with the new trend.

In a recent paper, Nelson and Startz (2007) suggest that UC models may suffer from weak identification and that in order to perform hypothesis testing, we should consider likelihood ratio (LR) test statistics.¹² Thus, we can compare models 1 and 2 in order to test for correlation. The LR test statistic for the null hypothesis that $\rho_{uv} = 0$ is 10.6, with p-value of 0.001. Therefore, we are able to reject the null of zero correlation at the 1% level of significance. As in Hryshko

¹² Nelson and Startz (2007) define weak identification as a situation in which the model is identified and asymptotic theory holds, but the data contain relatively little information about parameter estimates.

(2008), the data seems to favor the presence of negative correlation and allowing for it results in more volatile transitory shocks than when the zero-correlation restriction is imposed. Next, we compare models 1 and 3 in order to test for asymmetry. The LR test statistic is 26.18; however, notice that under the null of a single state, transition probabilities are not identified and therefore, this test is nonstandard and parametric bootstrapping is necessary in order to establish statistical significance – the results are pending, but with the large LR statistic we are confident that the null of symmetry will be reject once we have the results.¹³

[Insert Figure 2]

Figure 2 presents the estimated transitory components for models 1 and 3. The scale in both panels is natural logarithm times 100 so these can be interpreted as the percentage deviations from the trend. It is clear that in the symmetric representation, the transitory component is recurrent and small in amplitude, whereas in the asymmetric representation, transitory movements appear to be much more important, especially during recessions. Notice further that the series look as if they were uncorrelated, or weakly positively correlated, suggesting that imposing symmetry can seriously affect any attempts to identify income shocks.

[Insert Figure 3]

Figure 3 presents the filtered probabilities of asymmetric transitory shocks. The results suggest that these are important in order to capture almost half of the NBER-dated recessions. We acknowledge, however, that the NBER dating procedure draws on a large information set and that the methodology is largely subjective, and therefore agreement with the NBER dating is not a requirement for a valid decomposition into permanent and transitory components. Nonetheless, we find the comparison quite illustrative.

¹³ See Hansen (1992), Garcia (1998), and Di Sanzo (2007).

The results are in line with Kim and Murray (2002) and others, in that we find that not all recessions are alike but rather they differ in terms of the relative contribution of permanent and transitory shocks. Specifically, there are six out of ten recessions that do not appear to be characterized by these asymmetric shocks: 1960:2 – 1961:1, 1969:4 – 1970:4, 1980:1 – 1980:3, 1981:3 – 1982:4, 1990:3 – 1991:1, and 2001:1 – 2001:4. However, careful inspection reveals that the model could be capturing two of them: the 1960:2 – 1961:1 recession seems to be found a little earlier, whereas the 1980:1 – 1980:3 recession appears to be relatively close to the cutoff probably of 0.5. Finally, we note that three of the remaining four recessions that do not appear to be captured by a pluck have been previously classified in the literature as recessions characterized by slow recoveries.¹⁴ As such, our findings are consistent with the idea that these are not caused by a transitory shock.

We now refer to two or three non-recessionary plucks. The first one takes place in 1975:3, and it essentially capturing the effect of the 1975 tax rebate. Several authors have followed Blinder and Deaton (1985) and adjusted the data accordingly. Instead, we decided to let our model deal with such temporary phenomenon so it is actually a good thing that we are able to identify it. As of the pluck or plucks following the 1990 – 1991 recession, these can be attributed to measurement error problems in quarterly income data. In particular, Ludvigson and Michaelides (1998) explain the difficulties faced by the Bureau of Economic Analysis (BEA) in trying to account for the retroactive tax increase imposed in the 1993. In the end, the BEA assumed that all tax payments arrived in the second quarter of each of the relevant years, even though it is almost certain that a large portion of them arrived in the first quarter.

¹⁴ See Koenders and Rogerson (2005), Holmes and Silverstone (2006), and McKay and Reis (2008).

In sum, it should be evident that if one pretends to identify income shocks based on these types of decompositions, the underlying assumptions regarding asymmetry and correlation between innovations play a very important role as they lead to very different results.

IV. Testing for the PIH – Complete Information

We now proceed to test the implications of the PIH when labor income is described by the asymmetric UC-UR model estimated in the previous section. To test these implications, we consider a representative-agent version of the PIH, where the representative agent receives the aggregate income process. In addition, we assume a constant real interest rate, r , equal to 0.01. The results are summarized in Table 3.

[Insert Table 3]

Overall, we find conflicting evidence with the PIH. In particular, consumption seems to under react to permanent innovations while responding to transitory ones almost as outlined by the theory. Notice that when we include services expenditure in our measure of consumption, the MPC out of permanent shocks is lowered and the MPC out of transitory shocks becomes statistically significant. Moreover, the switching drift is significant in only one of our specifications. To complement the analysis, we also report an implied smoothness ratio- had the PIH hold – equal to 1.17; considerably higher than those in Table 1.

V. Testing for the PIH – Imperfect Information

In this section, we relax the assumption that agents can observe both components of income separately. Instead, we assume that agents understand the overall structure of the economy (i.e., they comprehend the process for aggregate income) but they can only observe the sum of components given in equation (2). However, we consider the case in which the relevant information becomes available with a one period lag. That is, in period t , agents observe a composite income shock, ε_t , as well as the realizations v_{t-1} , u_{t-1} and S_{t-1} , where ε_t equals:

$$\varepsilon_t = v_t + u_t + \pi S_t. \quad (12)$$

Notice that the composite income shock is generated by the process outlined in section II.A, but the consumer is unable to distinguish between permanent and transitory shocks. Therefore, we can think of the optimal consumption response as a combination of two parts: one in which the agent responds to the composite shock, and another in which she corrects for the error made in the previous period. We propose the following *ad-hoc* equation for the aggregate change in consumption:

$$\Delta \tilde{c}_t = a_0 + a_1 \cdot \varepsilon_t + a_2 \cdot \varepsilon_{t-1} + (a_3 \cdot v_{t-1} + a_4 \cdot u_{t-1} + a_5 \cdot \gamma_{S_{t-1}}). \quad (13)$$

Intuitively, when agents lack contemporaneous information, they are involved in a complicated signal extraction problem that results in a_1 – the MPC out of the composite shock.¹⁵ In addition, as information becomes available, the consumer may reverse what she did in the previous period, and consume as if she would have under the full information setting. Thus, we would expect $a_1 = -(1 + r)a_2$, and the terms in parentheses to match those implied by the PIH

¹⁵ In general, the consumer will attribute part of the composite shock to each of its components given their relative variances. However, in this context, the prediction for a_1 is quite complicated not only due to the Markov-switching nature of the asymmetric transitory shock but also because income innovations are correlated.

given by equation (8). As a result, if the intuition holds, consumption would appear to be excessively smooth to current income innovations and excessively sensitive to lagged income changes.

[Insert Table 4]

Table 4 presents the results from estimating equation (13). Notice that we do not restrict the coefficients on ε_t and ε_{t-1} to be equal to one another. Instead, we estimate the unrestricted model and we test the null hypothesis that $a_1 = -(1 + r)a_2$. Overall, we are unable to reject this null for any of our consumption measures – the p-values are 0.85, 0.51, and 0.70, respectively. This is consistent with the idea that agents revise their consumption decisions once information becomes available. As for the estimates of a_1 and a_2 , we are incapable of comparing them to the ones implied by the theory. However, given the relative variance of the different components of ε_t , it seems reasonable that they range from 0.3 to 0.5.

In terms of MPCs, we are able to reject the null that the consumption response to (lagged) permanent shocks equals one in two specifications. However, for nondurables expenditure, the p-value equals 0.22, which implies that we fail to reject the null that the MPC out of a permanent shock equals one. In turn, the MPC out of transitory shocks is only statistically insignificant for services, but we cannot reject the null that it is equal to the prediction of the PIH (i.e., that it is equal to 0.03) in the other specifications at the 10% and 1% level of significance, respectively. Finally, notice that the switching drift becomes statistically significant in our nondurables regression, and we cannot reject the null that it is equal to one with p-value of 0.12.

Therefore, if we focus on nondurables expenditure as our preferred measure of consumption, our intuitive imperfect information version of the PIH performs reasonably well. In particular, we find evidence supporting equation (13) not only in terms of MPCs but also regarding the

implied smoothness ratio, which equals 0.90 – less than 1.12 times higher than that observed in the actual data. Notice further that this ratio lies within two standard deviations from the actual statistic reported in Table 1. However, this version of the PIH fails when it comes to matching excess sensitivity. The OLS coefficient of a regression of the predicted growth rate of consumption on lagged income growth equals 0.77, which is about 3.5 times higher than that in the data.

A. *Testing for the PIH – Mixed Economy*

In the spirit of Campbell and Mankiw (1989), we consider an economy populated by agents that differ in their respective informational assumptions. In particular, we assume that there is a fraction δ of consumers that behave according to the PIH with full information, and that the remaining $(1-\delta)$ of consumers behave according to the lagged information framework.¹⁶ Therefore, we can characterize the change in aggregate consumption by the following expression:

$$\Delta \tilde{c}_t = \delta \cdot \Delta \tilde{c}_t^F + (1 - \delta) \cdot \Delta \tilde{c}_t^I, \quad (14)$$

where $\Delta \tilde{c}_t^F$ is the prediction of the PIH for the change in consumption under full information and $\Delta \tilde{c}_t^I$ is its incomplete information counterpart.

Rather than estimating equation (14), we look at the effect of δ on the implied variance of consumption growth and, in turn, on the implied smoothness ratio. That is, if we take the variance on both sides of equation (14) as follows:

¹⁶ Technically, δ is the fraction of income which flows to consumers who behave according to the full information version of the PIH.

$$Var(\Delta\tilde{c}_t) = \delta^2 \cdot Var(\Delta\tilde{c}_t^F) + (1 - \delta)^2 \cdot Var(\Delta\tilde{c}_t^I) + 2 \cdot \delta \cdot (1 - \delta) \cdot Cov(\Delta\tilde{c}_t^F, \Delta\tilde{c}_t^I), \quad (15)$$

and then we divide both sides of equation (15) by the variance of labor income growth, we can take the square root and solve for the implied smoothness ratio as a function of δ . Table 5 summarizes the relevant information entering equation (15).

[Insert Table 5]

In Figure 4 we plot the implied smoothness ratio (solid line) for different values of δ . It is worth noting that even though there is no value of δ that would reproduce the observed smoothness of consumption (long dashed line), we find that there is a wide range of values for which the smoothness ratio is below one – i.e., $\delta \in [0, 0.75]$ – and more importantly, a range of values for which such ratio lies within two standard deviations from the actual ratio for nondurables expenditure – i.e., $\delta \in [0, 0.65]$. In addition, we also report the implied sensitivity coefficient (short dashed line) as a function of δ .¹⁷ Importantly, we find that when $0.6 \leq \delta \leq 0.65$, the model is capable of matching the relative smoothness of consumption, while replicating the observed excess sensitivity.

[Insert Figure 4]

Therefore, our results suggest that we are able to match the consumption data if we consider an economy in which 60 – 65% of consumers behave according to the PIH with full information and the remaining consumers have delayed information

¹⁷ For a given value of δ , we generate the artificial series $\Delta\tilde{c}_t$ as in equation (14). Then, we regress $\Delta\tilde{c}_t$ on lagged income growth in order to calculate the sensitivity coefficient.

VI. Conclusion

In this paper, we decompose aggregate labor income into the sum of a permanent random walk component and a stationary transitory component. Contrary to most of the existing literature, we allow for asymmetric transitory movements and correlation between permanent and transitory innovations. Our results support the introduction of these features as it allows us to fit the data significantly better than common restrictive alternatives. Overall, this is the main contribution of the paper.

Next, we investigate whether the predictions of the Permanent Income Hypothesis conditional on this process for income can help explain the excess smoothness puzzle. The results suggest that under the assumption of full information, the model performs poorly: the marginal propensity to consume out of permanent shocks is significantly lower than one and consumption is too smooth relative to the predictions of the model. Therefore, we consider an extension in which consumers lack contemporaneous information but all information becomes available with one period lag. An intuitive version of the model in which agents respond partly to current income shocks and partly to newly arrived information on the composition of lagged income shocks fits the data significantly better. However, the correlation between consumption growth and lagged income growth is exaggerated.

Finally, we ask what would it take in order to be able to fully match the data. In this sense, we consider an economy populated by agents that differ in their respective informational assumptions. Our results show that a model in which 60 – 65% of consumers behave according to the PIH with full information while the remaining consumers have delayed information is capable of quantitatively reproducing the stylized facts of aggregate U.S. consumption data.

References

- Attanasio, O., and N. Pavoni (2007). "Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption." NBER Working Paper # 12994.
- Beaudry, P., and G. Koop (1993). "Do Recessions Permanently Change Output?" Journal of Monetary Economics **31**: 149-164.
- Blinder, A., and A. Deaton (1985). "The Time Series Consumption Function Revisited." Brookings Papers on Economic Activity **2**: 165-511.
- Campbell, J. Y., and N. G. Mankiw (1989). "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." NBER Macroeconomics Annual: 185-216.
- Carroll, C. D. (1992). "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." Brookings Papers on Economic Activity **2**: 165-511.
- Carroll, C. D., and M. Sommer (2003). "Epidemiological Expectations and Consumption Dynamics." Manuscript, Johns Hopkins University.
- Di Sanzo, S. (2007). "Testing for Linearity in Markov Switching Models: A Bootstrap Approach." Statistical Methods and Applications, **18**(2): 153-168.
- Flavin, M. A. (1981). "The Adjustment of Consumption to Changing Expectations About Future Income." Journal of Political Economy **89**(5): 974-1009.
- Friedman, M. (1969). Monetary Studies of the National Bureau. The Optimum Quantity of Money and Other Essays. M. Friedman. Chicago, Aldine: 261-284.
- Friedman, M. (1993). "The 'Plucking Model' of Business Fluctuations Revisited." Economic Enquiry **31**(2): 171-177.

- Garcia, R. (1998). "Asymptotic Null Distribution of the Likelihood Ratio Test in Markov Switching Models." International Economic Review **39**(3): 763-788.
- Goodfriend, M. (1992). "Information-Aggregation Bias." American Economic Review **82**(3): 508-519.
- Hall, R. E. (1978). "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." Journal of Political Economy **86**(6): 971-987.
- Hall, R. E., and F. S. Mishkin (1982). "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." Econometrica 50(2): 461-481.
- Hamilton, J. D. (1989). "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." Econometrica **57**(2):375-384.
- Hansen, B. E. (1992). "The Likelihood Ratio Test Under Non-Standard Conditions: Testing the Markov Switching Model of GNP." Journal of Applied Econometrics **7**: S61-S82.
- Hansen, L. P., and T. J. Sargent (1981). "A Note on Wiener-Kolmogorov Prediction Formulas for Rational Expectations Models." Economics Letters **8**(3): 255-260.
- Holmes, M. J., and B. Silverstone (2006). "Okun's Law, Asymmetries and Jobless Recoveries in the United States: A Markov-Switching Approach." Economics Letters **92**: 293-299.
- Hryshko, D. (2008). "Excess Smoothness of Consumption in an Estimated Life Cycle Model." Mimeo, University of Alberta.
- Kim, C.-J. (1994). "Dynamic Linear Models with Markov-Switching." Journal of Econometrics **60**: 1-22.

- Kim, C.-J., and C. J. Murray (2002). "Permanent and Transitory Components of Recessions." Empirical Econometrics **27**(2): 163-183.
- Kim, C.-J., and C. R. Nelson (1999). "Friedman's Plucking Model of Business Fluctuations: Tests and Estimates of Permanent and Transitory Components." Journal of Money, Credit, and Banking **31**(3): 317-334.
- Koenders, K., and R. Rogerson (2005). "Organizational Dynamics Over the Business Cycle: A View on Jobless Recoveries." Federal Reserve Bank of St. Louis Review **87**(4): 555-579.
- Ludvigson, S. C., and A. Michaelides (1998). "Can We Explain the 'Consumption Excesses' Yet? Aggregate Consumption Implications of Buffer Stock Saving Behavior." Working paper, Federal Reserve Bank of New York.
- Ludvigson, S. C., and A. Michaelides (2001). "Does Buffer-Stock Saving Explain the Smoothness and Excess Sensitivity of Consumption?" American Economic Review **91**(3): 631-647.
- MaCurdy, T. E. (1982). "The Use of Time Series Processes to Model the Error Structure of Earnings in Longitudinal Data Analysis." Journal of Econometrics **18**(1): 83-114.
- McKay, A., and R. Reis (2008). "The Brevity and Violence of Contractions and Expansions." Journal of Monetary Economics **55**: 738-751.
- Morley, J. C., C. R. Nelson, and E. Zivot (2003). "Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different?" The Review of Economics and Statistics **85**(2): 235-243.
- Nelson, C. R., and R. Startz (2007). "The Zero-Information-Limit Condition and Spurious Inference in Weakly Identified Models." Journal of Econometrics **138**(2007): 47-62.

- Pischke, J.-S. (1995). "Individual Income, Incomplete Information and Aggregate Consumption." Econometrica **63**(4): 805-840.
- Quah, D. (1990). "Permanent and Transitory Movements in Labor Income: An Explanation For 'Excess Smoothness' in Consumption." Journal of Political Economy **98**(3): 449-475.
- Sichel, D. E. (1994). "Inventories and the Three Phases of the Business Cycle." Journal of Business and Economic Statistics **12**(3): 269-277.
- Sinclair, T. M. (2010). "Asymmetry in the Business Cycle: Friedman's Plucking Model with Correlated Innovations." Studies in Nonlinear Dynamics and Econometrics **14**(1) article 3.
- Stock, J. H., and M. W. Watson (1988). "Variable Trends in Economic Time Series." Journal of Economic Perspectives **2**: 147-174.
- Watson, M. W. (1988). "Univariate Detrending Methods with Stochastic Trends." Journal of Monetary Economics **18**: 49-75.






Figure 1. Consumption Growth (demeaned) and NBER Recession Dates



Panel A - Symmetric UC-UR Model



Panel B - Asymmetric UC-UR Model

Figure 2. Estimated Transitory Components and NBER Recession Dates




Figure 3. Probabilities of Asymmetric Transitory Shocks




Figure 4. Actual vs. Implied Smoothness and Sensitivity Ratios

Note: The line labeled “Smoothness” reports the implied ratio of the standard deviation of the aggregate consumption growth to the standard deviation of aggregate labor income growth as a function of δ . The line labeled “Sensitivity” reports the OLS coefficient of consumption growth on lagged labor income growth as a function of δ . δ is defined as the fraction of consumers that behave according to the PIH with full information so that $(1-\delta)$ is the fraction of consumers that behave according to the lagged information framework. The line labeled “Actual Smoothness” reports the ratio of the standard deviation of aggregate consumption growth to the standard deviation of aggregate labor income growth for nondurables expenditure – see Table 1.

Table 1 – Relative Smoothness and Excess Sensitivity of Aggregate Consumption

	Relative Smoothness	Excess Sensitivity
ΔC_t	0.53 (0.04)	0.15 (0.03)
ΔC_t^{ND}	0.81 (0.06)	0.22 (0.05)
ΔC_t^S	0.51 (0.04)	0.09 (0.03)

Note: The column labeled “Relative Smoothness” reports the ratio of the standard deviation of the aggregate consumption growth measure in the row, to the standard deviation of aggregate labor income growth. In parentheses are the standard errors for this ratio, computed by GMM. The column labeled “Excess Sensitivity” reports the OLS coefficient of consumption growth on lagged labor income growth. OLS standard errors are in parentheses. The consumption growth measure ΔC_t is the growth in real, per capita nondurables and services expenditure. ΔC_t^{ND} is nondurables expenditure growth and ΔC_t^S is growth in services expenditure. Labor income is compiled from NIPA components as wages and salaries plus transfer payments, plus employer contributions for employee pensions and insurance, minus employee contributions for social insurance, minus taxes. This measure is also per capita and is deflated by the PCE chain-type price deflator. Sample 1952:1 – 2009:4.

Table 2 – Maximum Likelihood Estimates of Labor Income

Parameters		Asy. UC-UR (1)	Asy. UC-0 (2)	Sym. UC-UR (3)
Log Likelihood		-256.39	-261.69	-269.48
Std. Dev. Permanent Innovation	σ_v	0.98 (0.07)	0.72 (0.04)	0.85 (0.04)
Std. Dev. Transitory Innovation	σ_u	0.45 (0.08)	0.00 -	0.02 (0.01)
Correlation between innovations	ρ_{uv}	-1.00 -	Restricted to be zero	-1.00 -
Drift	μ	0.55 (0.07)	0.55 (0.05)	0.56 (0.06)
AR(1)	ϕ_1	0.14 (0.07)	0.16 (0.17)	1.69 (0.01)
AR(2)	ϕ_2	0.49 (0.07)	0.36 (0.19)	-0.99 (0.01)
Asymmetric shock	π	-2.21 (0.21)	-2.34 (0.31)	Restricted to be zero
Pr[St=0 St-1=0]	p_{00}	0.97 (0.01)	0.98 (0.01)	N/A
Pr[St=1 St-1=1]	p_{11}	0.58 (0.17)	0.61 (0.17)	N/A

Note: The column labeled “Asy. UC-UR” reports the parameter estimates of the asymmetric, unrestricted UC model presented section II.A. The column labeled “Asy. UC-0” reports the parameter estimates of the model when we impose the zero-correlation restriction. The column labeled “Sym. UC-UR” reports the parameter estimates of the model when we impose symmetry. Standard errors are in parentheses, except for ML estimates that fell on the boundary, which violates the regularity condition. To calculate the standard errors we treated these parameters as known constants for the purpose of calculating the second derivatives of the log likelihood. Labor income is compiled from NIPA components as wages and salaries plus transfer payments, plus employer contributions for employee pensions and insurance, minus employee contributions for social insurance, minus taxes. This measure is also per capita and is deflated by the PCE chain-type price deflator. Sample 1952:1 – 2009:4.

Table 3 – Testing for the PIH: Complete Information

Dependent Variable		ΔC_t	ΔC_t^{ND}	ΔC_t^S
Constant	-	0.57***	0.34***	0.58***
		(0.03)	(0.05)	(0.03)
Permanent Shock	v_t	0.20***	0.25***	0.17***
		(0.03)	(0.06)	(0.04)
Transitory Shock	u_t	0.12	0.09	0.17**
		(0.08)	(0.12)	(0.08)
Switching Drift	γ_{S_t}	2.01*	2.26	1.71
		(1.16)	(1.82)	(1.16)
Adjusted R-squared		0.16	0.12	0.11

Note: Model: $\Delta \tilde{c}_t = a_0 + a_1 v_t + a_2 u_t + a_3 \gamma_{S_t}$, where $\Delta \tilde{c}_t$ is the aggregate consumption growth rate measure in the column, v_t and u_t are permanent and transitory shocks identified by the “Asy. UC-UR” model estimated in Table 2, and γ_{S_t} is the switching drift term given in equation (9). OLS standard errors are in parentheses. The consumption growth measure ΔC_t is the growth in real, per capita nondurables and services expenditure. ΔC_t^{ND} is nondurables expenditure growth and ΔC_t^S is growth in services expenditure. Sample 1952:3 – 2009:4. ***, **, and * indicate significance at the 1, 5, and 10% level, respectively.

Table 4 – Testing for the PIH: Imperfect Information

Dependent Variable		ΔC_t	ΔC_t^{ND}	ΔC_t^S
Constant	-	0.59***	0.34***	0.60***
		(0.03)	(0.05)	(0.03)
Composite Shock	ε_t	0.26***	0.34***	0.20***
		(0.03)	(0.05)	(0.03)
Lagged Composite Shock	ε_{t-1}	-0.23*	-0.48**	-0.15
		(0.13)	(0.21)	(0.13)
Lagged Permanent Shock	v_{t-1}	0.41***	0.73***	0.25*
		(0.13)	(0.22)	(0.14)
Lagged Transitory Shock	u_{t-1}	0.26*	0.54**	0.10
		(0.14)	(0.24)	(0.15)
Lagged Switching Drift	$\gamma_{S_{t-1}}$	1.23	6.46*	-1.05
		(2.13)	(3.50)	(2.20)
Adjusted R-squared		0.32	0.25	0.20

Note: Model: $\Delta \tilde{c}_t = a_0 + a_1 \cdot \varepsilon_t + a_2 \cdot \varepsilon_{t-1} + (a_3 \cdot v_{t-1} + a_4 \cdot u_{t-1} + a_5 \cdot \gamma_{S_{t-1}})$, where $\Delta \tilde{c}_t$ is the aggregate consumption growth rate measure in the column, ε_t is the composite shock given in equation (12), v_t and u_t are permanent and transitory shocks identified by the “Asy. UC-UR” model estimated in Table 2, and γ_{S_t} is the switching drift term given in equation (9). OLS standard errors are in parentheses. The consumption growth measure ΔC_t is the growth in real, per capita nondurables and services expenditure. ΔC_t^{ND} is nondurables expenditure growth and ΔC_t^S is growth in services expenditure. Sample 1952:4 – 2009:4. ***, **, and * indicate significance at the 1, 5, and 10% level, respectively.

Table 5 – Covariance Matrix

	$\Delta\tilde{c}_t^F$	$\Delta\tilde{c}_t^I$	Δy_t
$\Delta\tilde{c}_t^F$	1.16	0.36	0.84
$\Delta\tilde{c}_t^I$	0.36	0.70	0.38
Δy_t	0.84	0.38	0.85

Note: The table presents the covariance between the variables entering equation (15). $\Delta\tilde{c}_t^F$ is the prediction of the PIH for the growth rate of consumption under full information, and $\Delta\tilde{c}_t^I$ is its imperfect information counterpart. Δy_t is the growth rate of aggregate labor income. Sample 1952:4 – 2009:4.