Macroeconomic Dynamics, **2**, 1998, 401–425. Printed in the United States of America.

TESTING THE PERMANENT-INCOME/LIFE-CYCLE HYPOTHESIS WITH AGGREGATE DATA

JOHN J. SEATER
North Carolina State University

The aggregate implications of the permanent-income/life-cycle hypothesis (PILCH) are derived rigorously. Virtually all empirical rejections of PILCH based on aggregated data are shown to result from misspecifications or from characteristics of aggregate data that have been overlooked. Valid aggregate tests are proposed. Those based on a properly formulated aggregate consumption function may be superior to those based on Euler-equation methods.

Keywords: Permanent-Income Hypothesis, Aggregate Data

1. INTRODUCTION

The permanent-income/life-cycle hypothesis (PILCH) is one of the greatest achievements of macroeconomic research. It offered the first coherent theory of consumption, reconciled seemingly contradictory implications of the data, and took the first steps toward providing macroeconomics with genuine microeconomic foundations. Development of intertemporal optimization methods has allowed a fully rigorous derivation of PILCH and its implications, and PILCH now is the theory that economists almost always use to analyze household choice. It is a tribute to the brilliance of the theory's founders that their original insights have survived the increase in rigor and generality with embellishment but no fundamental alteration.

Despite this theoretical success, empirical evidence on PILCH is contradictory. Microdata generally support PILCH, but aggregate data generally reject it (Deaton, 1992). The aggregate rejections are of several types: rejections of overidentifying restrictions in Euler-equation tests, the Deaton paradox, excess sensitivity of consumption to current income, statistical insignificance of the interest rate as an explanatory variable in consumption functions, age-consumption profile behavior, and predicted time-series behavior for consumption that is inconsistent with the data. These empirical results have had considerable impact on macroeconomic research, instigating a multitude of attempts to reconcile the evidence with

I thank David Weil for perceptive criticisms of an earlier version that prompted me to push to its logical conclusion the theory developed in Section 4. Address correspondence to: John J. Seater, Economics Department, Box 7506, North Carolina State University, Raleigh, NC 27695, USA; e-mail: john_seater@ncsu.edu.

 1365-1005/98 $9.50

PILCH and leading some to question PILCH's adequacy as a theory of household choice.

I argue herein that the reconciliation between the micro- and macroevidence on PILCH is simpler than has been recognized heretofore. Essentially, the macroevidence is largely invalid or at least suspect. PILCH is a theory of how atomistic agents behave, whereas aggregate data arise from the general-equilibrium behavior of the entire economy. The existing literature generally has proceeded by assuming that PILCH's aggregate implications are the same as its microimplications and that the tests appropriate for the microdata also are appropriate for aggregate data. Those assumptions often are incorrect because PILCH's atomistic and aggregate implications typically differ. Even when the implications are the same, the nature of aggregate data imposes special restrictions irrelevant to the microdata. The literature has not recognized these distinctions, leading to theoretical and empirical treatments that usually are invalid. The analysis presented below derives PILCH's proper aggregate implications and presents valid tests. The analysis suggests that, for testing PILCH, a properly formulated version of the traditional aggregate consumption function may provide a better foundation than the Euler-equation methods that have been in vogue in recent years.

2. ATOMISTIC AND REPRESENTATIVE AGENTS

The subsequent analysis hinges on a comparison of atomistic and representative agents. The discussion requires some general theory of those two agents' optimization problems, presented in this section. The theory is well known, and so, I omit details and derivations.[1]

2.1. Atomistic Agent

The atomistic agent maximizes his lifetime utility

$$\sum_{j=0}^{\infty}(1+\rho)^{-(j+1)}u(c_{t+j}) \tag{1}$$

subject to his lifetime budget constraint

$$\sum_{j=0}^{\infty}R_{t+j}c_{t+j} = A_t + \sum_{j=0}^{\infty}R_{t+j}y_{t+j}, \tag{2}$$

where ρ is the rate of time preference, u is the utility function, t is the initial time, c_{t+j} is current consumption, A_{t+j} is the stock of assets, y_{t+j} is current income, and R_{t+j} is the discount factor,

$$R_{t+j} = \prod_{k=0}^{j}(1+r_{t+k})^{-1},$$

where r_t is the real interest rate. I treat utility as a function of consumption only and income as exogenous to the individual. Nothing important depends on these restrictions. Solution is by the discrete-time maximum principle to facilitate comparison with the representative-agent solution, which requires a dynamic solution method.

The current-value Hamiltonian is

$$H_t = u(c_{t+j}) + \psi_{t+j+1}(r_{t+j}A_{t+j} + y_{t+j} - c_{t+j}) \tag{3}$$

and the necessary conditions are

$$A_{t+j+1} - A_{t+j} = r_{t+j}A_{t+j} + y_{t+j} - c_{t+j}, \tag{4}$$

$$\psi_{t+j+1} - \psi_{t+j} = (\rho - r_{t+j})(1+\rho)^{-1}\psi_{t+j+1}, \tag{5}$$

$$A_t, \quad \text{given} \tag{6}$$

$$\lim_{j-\infty} \psi_{t+j+1}(1+\rho)^{-(j+1)}A_{t+j} = 0, \tag{7}$$

$$u'(c_{t+j}) = \psi_{t+j+1}. \tag{8}$$

Equation (7) is the transversality condition; equations (4), (6), and (7) together are equivalent to the original lifetime budget constraint (2). I use the necessary conditions later to construct Euler-equation tests and the consumption function for the atomistic individual.

2.2. Representative Agent

Aggregate data are generated by the general-equilibrium behavior of the economy. One way to model aggregate behavior is to solve for the competitive equilibrium among the economy's atomistic agents. That approach almost always is analytically intractable. A tractable approach is to recast the general-equilibrium problem as a representative-agent problem. The First Fundamental Theorem of Welfare Economics establishes that a decentralized general-equilibrium solution can be achieved as the solution to a planning problem. Under well-known restrictions on individual utility functions, a representative agent exists, and the planning problem in question is equivalent to the representative agent's utility maximization problem. There are several possible variants on the necessary restrictions, but for the issues discussed below, it makes no difference which set of restrictions is imposed.[2]

The restrictions required for existence of a representative agent are very severe and hardly anyone believes that they really hold. The issue is whether they produce a model (the representative-agent model) that is a good approximation to the general-equilibrium behavior of the economy. Many critics [e.g., Kirman (1992)] do not believe so and argue that PILCH fails to explain the aggregate data simply because there is no representative agent to whom PILCH applies, even if PILCH does explain how every individual agent behaves. I have no desire to enter this debate here; rather, I will accept the representative agent as equivalent to the solution of competitive equilibrium. My point is that, even if the representative-agent framework is valid, there are important differences between PILCH's implications for atomistic and representative agents.

The representative agent's problem is similar but not identical to that facing an atomistic agent; the important difference between the two problems is the budget constraint. The representative agent's lifetime utility function looks just like that of the atomistic agent,

$$\sum_{j=0}^{\infty}(1+\rho)^{-(j+1)}u(c_{t+j}),$$

except that c is now consumption per capita and the utility function u need not be the same as that of any individual agent. The representative agent maximizes his lifetime utility subject to his budget constraint, which is that consumption plus gross investment equals output:

$$k_{t+j+1} - k_{t+j} = f(k_{t+j}; \Omega_{t+j}) - c_{t+j} - \delta k_{t+j}.$$

This equation, of course, is the law of motion for the capital stock in a closed economy. (I discuss the open-economy case later.) Here, f is the production function, Ω_t is the state of technology at time t, and δ is the depreciation rate. The representative agent also is subject to (1) the initial condition that the initial capital stock equal its given value k_t; (2) a transversality condition, discussed momentarily; and (3) the control constraint,

$$0 \le c_{t+j} \le f(k_{t+j}; \Omega_{t+j}). \qquad \textbf{(9)}$$

The lower bound on current consumption in (9) is irrelevant if $u'(c) \to \infty$ as $c \to 0$, as under the frequently imposed Inada conditions. The upper bound, however, always is relevant as long as $u'(c) > 0$ for all values of c.

The upper bound $f(k; \Omega)$ in (9) changes over time as k changes. The representative agent himself determines the path of k as part of his optimal plan. It is this endogenous dynamic element of the budget constraint that requires us to use a dynamic method to solve the representative agent's problem. The atomistic agent does not face a similar situation. His consumption choice in any particular time period is unaffected by the current value of his state variable (his asset holdings) because he is free to borrow or lend as much as he chooses in any given period. His only constraint is that at the end of his life he must hold no assets, positive or negative. In contrast, the representative agent is equivalent to the entire economy and thus has no one to borrow from. Consequently, he cannot break the tight connection between the current value of his state variable (the capital stock) and his current and future opportunities for consumption and investment. This difference in constraints between the atomistic and representative agents is central to the discussion in Section 4 and is explored further there.

The representative agent's current-value Hamiltonian is

$$H_{t+j} = u(c_{t+j}) + \lambda_{t+j+1}[f(k_{t+j};\Omega_{t+j}) - c_{t+j} - \delta k_{t+j}]$$

where λ is the costate variable. The necessary conditions are

$$k_{t+j+1} - k_{t+j} = f(k_{t+j};\Omega_{t+j}) - c_{t+j} - \delta k_{t+j}, \tag{10}$$

$$\lambda_{t+j+1} - \lambda_{t+j} = [\rho + \delta - f'(k_{t+j};\Omega_{t+j})](1+\rho)^{-1}\lambda_{t+j+1}, \tag{11}$$

$$k_t, \quad \text{given} \tag{12}$$

$$\lim_{j - \infty} \lambda_{t+j+1}(1+\rho)^{-(j+1)}k_{t+j} = 0, \tag{13}$$

$$u'(c_{t+j}) = \lambda_{t+j+1} \quad \text{for all} \quad j, \tag{14}$$

and also (9). Equation (13) is the transversality condition and is a kind of terminal condition on the state variable k. As with the atomistic agent, we use the necessary conditions to construct Euler-equation tests and the consumption function for the representative agent.

3. EULER-EQUATION TESTS

The favored tests of PILCH over the past 15 years or so have been based on Euler-equation methods, and so, I begin with those. The nature of aggregate data creates serious difficulties for the Euler-equation approach.

3.1. Nature of Euler-Equation Tests

Euler-equation tests are based on estimating the Euler equation,

$$\frac{u'(c_{t+i+1})}{u'(c_{t+i})}\frac{1+r_{t+i+1}}{1+\rho} = 1, \tag{15}$$

obtained by taking the ratio of (8) in adjacent time periods. Under the conditions discussed in Section 2. 2, we can use (14) instead of (8) to do the same thing for the representative agent, obtaining (15) again. The Euler equation thus is the same for atomistic and representative agents. When we introduce uncertainty, (15) becomes

$$E_{t+i+1}\left[\frac{u'(c_{t+i+1})}{u'(c_{t+i})}\frac{1+r_{t+i+1}}{1+\rho}\right] = 1$$

or

$$\frac{u'(c_{t+i+1})}{u'(c_{t+i})}\frac{1+r_{t+i+1}}{1+\rho} = 1 + \varepsilon_{t+i+1}, \tag{16}$$

where ε is a mean-zero error.

Equation (16) implies that expected consumption in any period depends only on that period's interest rate and the previous period's consumption. The Euler-equation test consists of expanding (16) to include other supposedly irrelevant

https://doi.org/10.1017/S1365100598008062 Published online by Cambridge University Press

variables and testing their statistical significance. To conduct the test, we must specify a utility function and an expectations mechanism. For discussion, I assume constant relative risk aversion, and so, the Euler equation is

$$\left(\frac{c_t}{c_{t-1}}\right)^{\alpha}\frac{1+r_t}{1+\rho} = 1+\varepsilon_t, \tag{17}$$

where $\alpha = \gamma - 1$. Equation (17) can be estimated and tested with generalized method of moments (GMM). Estimation requires one moment condition for each parameter of the model; testing requires at least one more than that. Equation (17) has two parameters, ρ and α, and so, estimation requires two moment conditions. The Euler equation itself is one, but it is the only one that PILCH provides. We need another just to estimate the model and at least one more beyond that to provide an overidentifying restriction for conducting a test. We can obtain those extra moment conditions by specifying an expectations mechanism. The literature always has assumed the strong form of rational expectations, supposing that agents form their expectations as conditional forecasts based on all available information. To obtain the required extra moment conditions, rewrite (17) as

$$E_{t-1}\left[\left(\frac{c_t}{c_{t-1}}\right)^{\alpha}\frac{(1+r_t)}{(1+\rho)} - 1\right] = 0,$$

where E_{t-1} denotes expectation conditional on the information set Φ_{t-1}. Then, for any variable z_{t-1} in Φ_{t-1}, we have

$$E_{t-1}\left\{\left[\left(\frac{c_t}{c_{t-1}}\right)^{\alpha}\frac{(1+r_t)}{(1+\rho)} - 1\right]z_{t-1}\right\} = z_{t-1}E_{t-1}\left[\left(\frac{c_t}{c_{t-1}}\right)^{\alpha}\frac{(1+r_t)}{(1+\rho)} - 1\right] = 0. \tag{18}$$

By assumption, any variable dated $t-1$ or earlier is in Φ_{t-1}, and so, we have available as many extra moment conditions as we want, simply by listing as many lagged variables as we want (the "instruments" in GMM jargon).

With aggregate data, (16) usually rejects PILCH; various lagged instruments typically have significant coefficients [e.g., Hansen and Singleton (1983) and Campbell and Mankiw (1990)].

3.2. Exact Aggregation vs. GARP

Euler-equation tests of PILCH amount to estimating the utility function and checking its consistency with the data. When we use aggregate data, we are attempting to estimate an aggregate utility function. Such a function may not exist, and the estimation exercise then is pointless. It often is argued that existence of an aggregate utility function is guaranteed if individual utility functions meet the requirements to make exact aggregation possible—conditions that are stringent and unlikely to be met in practice. In fact, however, exact aggregation is neither necessary

https://doi.org/10.1017/S1365100598008062 Published online by Cambridge University Press

nor sufficient for Euler-equation estimation with aggregate data to be valid. What matters is whether the data are consistent with the generalized axiom of revealed preference (GARP).

Euler-equation estimation requires existence of a utility function capable of generating the data. Aggregate data are time-series data, and so, we need a utility function that is constant over time. Exact aggregation requires a restriction on the form of individual utility functions (that the indirect utility function be of Gorman form), but does not require constancy of those functions over time. As long as the utility functions remain of the correct form, exact aggregation is possible. Changing individual functions, however, imply a changing aggregate utility function. Estimating an Euler equation with aggregate data will be impossible in such a case, even though exact aggregation is possible in every period. Exact aggregation thus is not sufficient for estimating Euler equations with aggregate data.

Exact aggregation is not a necessary condition, either. What is necessary is existence of a utility function consistent with the data. The function need not result from exact aggregation. Varian (1982, 1983) proves that a utility function consistent with a data set exists if and only if the data are consistent with GARP, and he provides tests of whether such consistency obtains. The tests are nonparametric and require no assumptions about the utility function's form. If the data pass the GARP tests, Euler-equation estimation is possible; otherwise, it is not because there is no utility function to estimate. Exact aggregation is irrelevant.

Fleissig et al. (1997) apply Varian's tests to the aggregate consumption data for the United States. They find that the data generally are not consistent with GARP over the full sample periods of data availability but are consistent over subsamples of substantial length. For those subsamples, additional tests developed by Varian show that the aggregate utility function does not have the separability structure usually assumed in the literature. Fleissig et al. (1997) apply these results to Euler-equation estimation. They find that when sample periods are restricted to those consistent with GARP and the proper separability structure is imposed on the utility function, the overidentifying restrictions are not rejected for the instruments they examine and the form and parameter magnitudes of the utility function differ from those obtained with GARP-inconsistent data. These results suggest that GARP inconsistency may explain the tendency of reported Euler-equation tests to reject PILCH when aggregate data are used. Further work is needed to explore the robustness of these results.

Although exact aggregation is irrelevant to the validity of Euler-equation estimation, it is very relevant to the interpretation of any estimates obtained. Suppose one estimates an Euler equation with aggregate data over a GARP-consistent subsample. How does one interpret the resulting parameters? If the conditions of exact aggregation are met, there is no problem; the aggregate utility function is closely related to the underlying atomistic utility functions. If the aggregation conditions are not met, however, estimated aggregate utility function generally bears no relation at all to the underlying atomistic utility functions. It then is difficult to see what interpretation to give the estimated parameter values.[3]

https://doi.org/10.1017/S1365100598008062 Published online by Cambridge University Press

3.3. Aggregator Functions

A major reason that Euler-equation tests have gained favor over the past 15 years is that they supposedly estimate deep structural parameters that are invariant to regime changes, whereas decision rules such as the consumption function vary with regime changes. This argument originated with Lucas (1976). Lucas's argument has been taken as a rationale for preferring Euler-equation methods to estimating decision rules, but in fact, if regime changes are important enough to invalidate estimation of decision rules, they also are likely to invalidate Euler-equation estimation. The problem is that the aggregator function used to construct the aggregate data is no more structural than are aggregate decision rules such as the consumption function (Geweke, 1985). Treating the aggregator as constant yields a set of misconstructed data if significant structural changes in the economy occur. Ignoring the dependence of aggregators on the underlying regime is no more or less appealing than ignoring the dependence of decision rules on the regime. Either both dependencies are important or neither is. If they are important, then the available aggregate data, such as NIPA, are at least as inappropriate for Euler-equation estimation as they are for structural estimation, and estimates obtained by applying Euler-equation methods to those data are unreliable.

3.4. Rational Expectations

Euler-equation estimation requires that one maintain as a joint hypothesis the strongest form of rational expectations (RE), according to which agents are aware of all available information and use it in a statistically optimal way to update their expectations. It is the strong RE hypothesis that makes the Euler-equation approach feasible by providing all but one of the moment conditions necessary for estimation and testing. This joint hypothesis is imposed quite routinely to obtain moment conditions, but there are strong reasons to doubt its validity.

There is much evidence that expectations often do not conform to the RE hypothesis. The RE hypothesis is simply an application of optimal statistical decision theory and, as such, rests on Bayes' rule (Meinhold and Singpurwalla, 1983). However, most people do not understand Bayes' rule or use it in making decisions (Salop, 1987), and they usually make *systematic* mistakes in estimating or revising probabilities (Machina, 1987). Moreover, Lovell (1986), Jeong and Maddala (1991), and Zarnowitz (1992, Ch. 16) present a wide range of direct evidence that observed expectations of many agents do not conform to the RE hypothesis. Lovell also reports evidence favorable to several alternative expectations schemes that would not imply (18). Under such circumstances, testing PILCH by Euler-equation methods would lead to rejections of the moment conditions, but it would be the RE hypothesis rather than PILCH that was the source of the rejection.

It seems especially unlikely that RE is an appropriate model for the representative agent's expectations concerning aggregate data. When PILCH is tested with microdata, the assumption motivating (18) is that individual agents have freely available the data on all lagged values of personal economic variables, such as income, and fully understand the information content of those data. It is not too serious a stretch of the imagination to accept this assumption as a good approximation and to believe further that the econometrician has less information about an individual's personal economic variables than the individual himself has. When PILCH is tested with aggregate data, however, the assumption motivating (18) is that people have freely available data on all lagged values of macroeconomic variables and fully understand their information content. This assumption is unbelievable. Data are not the same as information. Most people do not understand the full meaning of data on GNP, inflation, and the like; nor do they understand how such concepts fit into a general-equilibrium economic model to provide predictive value, as anyone who has taught economics knows.[4] The data may be free, but their information content is not. Given that fluctuations in an individual's income are much more influenced by personal factors than by aggregate behavior (Pischke, 1995), the average person is unlikely to find it worthwhile to bear the costs of collecting and learning how to interpret aggregate data (Feige and Pearce, 1976). In that case, even though data on lagged macrovariables would help predict aggregate behavior, individuals would not use them in making decisions. It therefore is likely that the macroeconometrician has *more* information about the behavior of the economy than does the typical individual, and we should expect to see individuals (and therefore the representative agent) not using all the information available in the data. We thus also should expect (18) to be violated when confronted with aggregate data, even if PILCH is true.[5]

4. AGGREGATE CONSUMPTION FUNCTION

The foregoing arguments suggest that testing PILCH by applying Euler-equation methods to aggregate data may not be informative. Further consideration suggests that the older approach of estimating and testing the consumption function may be superior. However, a careful examination of PILCH's implications for aggregate consumption shows that the usual tests reported in the literature are misspecified, and several famous rejections of PILCH, such as excess sensitivity of consumption to current income and the Deaton paradox, are either consistent with PILCH or simply uninformative of its validity. Valid tests exist, but they await application to the data.

4.1. Consumption Functions vs. Euler Equations

In recent years, tests of PILCH based on consumption-function estimation have been regarded as inferior to Euler-equation tests for at least two reasons. One reason is that, as already mentioned, Lucas's (1976) critique argued that consumption functions are not stable with respect to regime changes; the other is that consumption-function estimation requires a joint hypothesis about the income-generating function to allow construction of a lifetime wealth or permanent-income

https://doi.org/10.1017/S1365100598008062 Published online by Cambridge University Press

series. Euler-equation tests do not suffer from these problems because utility functions are invariant to regime changes and no joint hypothesis on income generation is needed to estimate them. However, these arguments are overstated.

First, there is no convincing evidence that significant regime changes have occurred within the time span of data typically used for estimation, and so, it is not obvious that consumption-function estimation suffers from instability problems due to regime changes.[6] Also, Lucas's critique was directed against the traditional methods of analyzing government policy interventions. Tests of PILCH do not involve policy interventions, and so, Lucas's critique is largely beside the point. Finally, Cooley et al. (1984) have shown that, in any case, regime changes do not invalidate estimation of proper general decision rules. Regime changes themselves are something that the rational agent knows can happen and so can incorporate as possibilities in his expectations function. (Think of a Baysian expectations rule that, upon confronting a new observation, updates the perceived distribution of the policy choices within each possible policy regime and also updates the distribution of regimes, leading to a new metadistribution of all possible policy choices.) The important implication of Lucas's critique is that decision rules that are too narrowly specified will be inaccurate, not that it is impossible to specify decision rules at all.

Second, although consumption-function estimation requires a joint hypothesis on the income process, Euler-equation tests require *two* joint hypotheses, the first of which is totally unnecessary to consumption-function estimation and the second of which is much stronger than the hypothesis maintained in consumption-function estimation. The first hypothesis is the form of the utility function, which Euler-equation estimation requires that one specify exactly. Misspecification of the utility function can invalidate the entire analysis, and indeed the frequent rejections of the overidentifying restrictions obtained in the literature has led to a search for the correct specification. Quadratic [e.g., Hall (1978)], constant relative risk aversion [e.g., Hansen and Singleton (1983)], *S*-branch [e.g., Eichenbaum and Hansen (1990)], and seminonparametric [e.g., Fleissig et al. (1997)] are just some of the forms that have been explored. The traditional consumption-function tests have an advantage in this regard. Most implications of PILCH for consumption behavior are independent of the form of the utility function, and one can test them without specifying the consumption function exactly. For example, PILCH predicts that the atomistic agent's consumption will be positively related to his lifetime wealth and unrelated to his current income. This prediction can be tested by estimating an approximation to the consumption function, such as a linear or log-linear relation. No joint hypothesis on the utility function's form need be maintained, in contrast to Euler-equation estimation.[7] Given that we cannot observe utility functions and have little direct knowledge of their characteristics, this advantage seems considerable. The second joint hypothesis maintained by Euler-equation estimation is the strong form of rational expectations. As discussed earlier, this hypothesis is hard to accept for the representative agent (equivalently, for the aggregate data). Again, consumption-function tests of PILCH have an advantage in this regard—two advantages, in fact. First, they require a joint hypothesis of some model of lifetime wealth, which amounts to a model of expectations of future income, but they do not require that those expectations be rational in the sense of the RE hypothesis. Second, the joint hypothesis in consumption tests concerns only future income, whereas the joint hypothesis in Euler-equation tests is strong RE with respect to everything. It seems at least as reasonable to impose on income alone a model that can be estimated from market data as it does to impose on all variables a universal model of expectations, whose behavior we usually do not observe at all.

Consumption-function tests impose fewer and weaker joint hypotheses than Euler-equation tests and are not obviously inferior to them. Examining PILCH's implications for aggregate consumption therefore is worthwhile. The results are surprising.

4.2. Atomistic Consumption

We need to compare atomistic and aggregate consumption, and so, we begin by deriving the two consumption functions.

The atomistic agent's consumption c_{t+j} at time $t + j$ is obtained by inverting (8) to obtain

$$c_{t+j} = \hat{c}(\psi_{t+j+1}), \qquad \hat{c}' < 0. \tag{19}$$

From (5), we obtain[8]

$$\psi_{t+j+1} = \psi_t (1 + \rho)^{j+1} R_{t+j}. \tag{20}$$

The initial value ψ_t of the costate variable is chosen to satisfy (7) or, equivalently, (2). By using (4), (19), and (20), we can rewrite (2) as

$$\sum_{j=0}^{\infty} R_{t+j} \hat{c}\left[\psi_t (1 + \rho)^{j+1} R_{t+j}\right] = A_t + \sum_{j=0}^{\infty} R_{t+j} y_{t+j} \equiv W_t, \tag{21}$$

which implies that ψ_t is a function of lifetime wealth, W_t, and the entire sequence of current and future interest rates,

$$\psi_t = \psi\left(W_{t'} \{r_{t+j}\}_{j=0}^{\infty}\right). \tag{22}$$

We then can substitute into (19) to obtain

$$c_t = c\left(W_{t'} \{r_{t+j}\}_{j=0}^{\infty}\right). \tag{23}$$

We have the derivatives

$$\frac{\partial c_t}{\partial W_t} > 0$$

$$\frac{\partial c_t}{\partial r_{t+j}} \gtreqless 0 \quad \text{for all} \quad j.$$

https://doi.org/10.1017/S1365100598008062 Published online by Cambridge University Press

The ambiguity in the effect of interest rates reflects opposing income and substitution effects. In general, c depends on the entire path of interest rates independently of that path's effect on W_t, reflecting intertemporal income and substitution effects.

The most important implication of (23) is the well-known PILCH result that atomistic consumption is independent of both current income and current assets and depends only on the value of lifetime wealth (which determines the initial value of ψ) and the relation between r and ρ. Saving is a residual that offsets fluctuations in current income and permits optimal smoothing of consumption.

4.3. Aggregate Consumption

Aggregate consumption is equivalent to the representative agent's consumption, c_{t+j}, which is obtained in a manner similar to that for the atomistic agent. Invert (14) to obtain

$$c_{t+j} = \hat{c}(\lambda_{t+j+1}), \qquad \hat{c}' < 0 \tag{24}$$

From (11),

$$\lambda_{t+j+1} = \lambda_t (1+\rho)^{j+1} \prod_{i=0}^{j} [1 + f'(k_{t+i}; \Omega_{t+i}) - \delta]^{-1}. \tag{25}$$

(Recall that Ω_{t+i} is the state of technology at time $t+i$.) The initial value λ_t of the costate variable is chosen to satisfy (13). Using (10), (11), and (24), we can rewrite (13) as

$$\begin{aligned} 0 = \lim_{j-\infty} \left\{ \lambda_t \prod_{i=0}^{j} [1 + f'(k_{t+i}; \Omega_{t+i}) - \delta]^{-1} \right\} \Bigg(\sum_{h=0}^{j} (1-\delta)^{j-h} f(k_{t+h}; \Omega_{t+h}) \\ - \sum_{h=0}^{j} \left\{ (1-\delta)^{j-h} \hat{c} \left[\lambda_t (1+\rho)^{h+1} \prod_{m=0}^{h} [1 + f'(k_{t+m}; \Omega_{t+m}) - \delta]^{-1} \right] \right\} \\ + (1-\delta)^{j+1} k_t \Bigg), \end{aligned} \tag{26}$$

which implies that λ_t is a function of initial capital k and the entire sequence of products and marginal products of capital,

$$\lambda_t = \hat{\lambda}\big(\{f(k_{t+j}; \Omega_{t+j})\}_{j=0}^{\infty}, \{f'(k_{t+j}; \Omega_{t+j})\}_{j=0}^{\infty}, k_t\big). \tag{27}$$

We then can substitute into (24) to obtain

$$c_t = c^*\big(\{f(k_{t+j}; \Omega_{t+j})\}_{j=0}^{\infty}, \{f'(k_{t+j}; \Omega_{t+j})\}_{j=0}^{\infty}, k_t\big). \tag{28}$$

Finally, we can repeatedly iterate (10) backward to replace each k_{t+j} with a function of initial k and the sequence $\{\Omega_{t+i}\}_{i=0}^{j}$ and then substitute into (27) to get

$$\lambda_t = \lambda\left(\{\Omega_{t+j}\}_{j=0}^{\infty}, k_t\right).$$

Substituting these results in (28) gives

$$c_t = \tilde{c}\left(\{\Omega_{t+j}\}_{j=0}^{\infty}, k_t\right). \quad \textbf{(29)}$$

We thus have the standard result that aggregate consumption depends on tastes, technology, and initial conditions.

We usually think of the aggregate consumption function as a relation between consumption on the one hand and income and interest rates on the other. We can derive such a relation easily enough. Simply substitute y and r for $f(k)$ and $f'(k)$, respectively, wherever the latter two expressions appear in (28) to obtain

$$c_t = c\left(\{y_{t+j}\}_{j=0}^{\infty}, \{r_{t+j}\}_{j=0}^{\infty}, k_t\right). \quad \textbf{(30)}$$

It can be shown that

$$\frac{\partial c_t}{\partial y_{t+j}} > 0 \quad \text{for all} \quad j, \quad \textbf{(31)}$$

$$\frac{\partial c_t}{\partial r_{t+j}} \begin{cases} \gtreqless 0 & \text{for} \quad j = 0 \\ < 0 & \text{for} \quad j > 0 \end{cases}, \quad \textbf{(32)}$$

$$\frac{\partial c_t}{\partial k_t} > 0, \quad \textbf{(33)}$$

and that

1. c_t depends on the entire sequence of current outputs and current interest rates from period t to the end of the planning horizon, and
2. the present value of income is absent from (30), and so, c_t does not depend on it.

The relation between aggregate consumption on the one hand and current income, wealth, and interest rates on the other is strikingly different from that for atomistic consumption. First, recall that the "independent" variables in (30) are not truly independent but rather are determined simultaneously with the "dependent" variable—aggregate consumption. Aggregate consumption depends only on tastes, technology, and initial conditions, as shown in (29). The functional relation in (30) is correlational, not causal and we must not think of aggregate income or the interest rate as "causing" aggregate consumption. Second, the list of independent variables in (30) differs from that in the atomistic consumption function (23). Atomistic consumption depends on wealth but not on the particular path of current income. In contrast, Property 1 tells us that aggregate consumption "depends on" the sequence of future incomes with equation (31) showing the sign of that dependence; Property 2 tells us that consumption does not depend on the present value of lifetime income, i.e., does not depend on lifetime wealth; and Properties 1 and 2 together then imply that aggregate consumption depends on the particular path of lifetime income. These results are exactly the opposite of those for atomistic consumption!

The differences between atomistic and aggregate consumption in their relation to current income and wealth arise from the different budget constraints facing the atomistic and representative agents. The atomistic agent's constraint is a static one requiring only that terminal wealth be zero. The particular values of wealth between the initial and terminal periods are irrelevant. It is precisely that irrelevance that allows the atomistic agent to borrow and lend freely to smooth his consumption irrespective of what current income happens to be. The representative agent has no one to borrow from or lend to, and so, his budget constraint is a dynamic one requiring that his consumption always be less than or equal to his current income, which is aggregate output. However, the representative agent has the power to change future values of the upper bound in his budget constraint by changing the amount of his current income that he devotes to investment, and so, he is not indifferent to a rearrangement of his current income path. A change in his current income changes in the same direction his ability to alter the bound in his budget constraint, and he will alter his consumption path in response to any such change in current income.

Similarly, interest rates can have both income and substitution effects on the atomistic agent but can have only substitution effects on the representative agent because there is no net financial wealth in the aggregate. That is why the response of atomistic consumption to an interest-rate change generally is ambiguous, whereas the response of aggregate consumption generally is not.

In analyzing aggregate behavior, one might be tempted to apply the First Welfare Theorem in a casual manner and conclude that the representative agent's problem is the same as the atomistic agent's. That conclusion is incorrect and results from a misunderstanding of what the First Welfare Theorem says. The theorem guarantees that, under the usual restrictions, one can analyze the competitive equilibrium of the economy as if it were the outcome of a suitable individual agent. The suitable agent, however, is not a true atomistic agent but rather is the representative agent, who has a fundamentally different budget constraint from that facing a true atomistic agent. As a result, aggregate consumption's behavior is not the same as atomistic consumption's behavior, as the foregoing derivations have shown.

4.4. Empirical Implications

Equation (30) is nonlinear, but as a first pass, we could estimate the linear approximation

$$c_t = a + bk_t + \sum_{i=0}^{I} d_i y_{t+i} + \sum_{j=0}^{J} e_j r_{t+j} + fW_t + \varepsilon_t \qquad \mathbf{(34)}$$

where I and J are some empirically determined truncation limits and W is lifetime wealth, defined in (21). (One can substitute permanent income for lifetime wealth without changing anything that follows. In some subsequent sections, it will be convenient to make this substitution.) The foregoing theory predicts $b > 0,\ d_i > 0\ \forall\ i,\ e_0 \gtrless 0,\ e_i < 0\ \forall\ i > 0$, and $f = 0$. These implications distinguish aggregate PILCH consumption from

1. PILCH atomistic consumption, for which $b = 0,\ d_i = 0\ \forall\ i,\ e_i \gtrless 0\ \forall\ i$, and $f > 0$;
2. Keynesian aggregate consumption, for which $b = 0,\ d_0 > 0,\ d_i = 0\ \forall\ i > 0,\ e_0 < 0,\ e_i = 0\ \forall\ i > 0$, and $f = 0$; and
3. rule-of-thumb aggregate consumption, for which $b = 0,\ d_0 = 1,\ d_i = 0\ \forall\ i > 0,\ e_i = 0\ \forall\ i$, and $f = 0$.[9]

Aggregate data thus can be used to test PILCH against the leading alternatives.

Equation (34) apparently has never been used to test PILCH. What has been used instead is

$$c_t = a + dy_t + er_t + fW_t + \varepsilon_t, \qquad (35)$$

which is almost the atomistic agent's consumption function, except that future interest rates have been omitted and, for the purpose of testing, current income has been included. When (35) is estimated with microdata, PILCH implies $d = 0,\ e \gtrless 0$, and $f > 0$. Those same restrictions also have been tested with aggregate data. Unfortunately, when treated as an aggregate consumption function, equation (35) is a misspecification of the true function (34) that omits several relevant variables (future incomes and interest rates) and includes an irrelevant variable (lifetime wealth). This misspecification causes severe econometric difficulties. Lifetime wealth is, by construction, correlated with all future incomes and interest rates and so may be significant in (35) by proxying for those omitted variables. Lifetime wealth also is correlated with current income and the current interest rate and may tend to rob those variables of significance. Finally, current income and the current interest rate are correlated with future incomes and interest rates because the dynamic adjustment path of the economy creates serial correlation in incomes and interest rates. In summary, the omitted variables are correlated with all of the included variables, the latter are correlated with each other, and the error term is serially correlated. The coefficient estimates obtained from equation (35) therefore will be biased, inconsistent, and inefficient. Their values will not indicate the true effects of current income, the interest rate, and lifetime wealth on aggregate consumption, and their standard errors will be misstated.[10] Obviously, estimates obtained from (35) are worthless and tell us nothing about the validity of PILCH.

4.5. Excess Sensitivity and Interest-Rate Insignificance

As we have seen, an atomistic individual's consumption path depends on his lifetime wealth but not the particulars of his current income path. PILCH therefore predicts that current income will be statistically insignificant in a regression of current consumption on lifetime wealth and current income. However, tests using aggregate data virtually always reject this implication by finding a significant effect of current income. This is the well-known problem of "excess sensitivity" of consumption to current income. Another apparent empirical problem with PILCH is insignificant sensitivity of aggregate consumption to the interest rate. The theory for the atomistic individual predicts a negative relation between the two, but the aggregate data routinely produce an insignificant relation with no consistent sign. Aggregate consumption thus seems excessively sensitive to current income and insufficiently sensitive to interest rates to be consistent with PILCH.

The results of the previous subsection suggest a radically different interpretation of these famous empirical rejections of PILCH, which is that they do not constitute rejections at all because they do not test true implications of PILCH and are econometrically unreliable. As we have seen, estimating equation (35) with aggregate data leads to results that are biased, inconsistent, and inefficient. Evidence in (35) of "excess sensitivity" of aggregate consumption to current aggregate income and insufficient sensitivity to interest rates does not reject PILCH; it simply is uninformative.

4.6. Deaton Paradox

Suppose a shock raises the present value of an atomistic individual's income path. The individual's optimal response to the income shock is to raise current consumption by approximately the amount that the shock raises his permanent income (i.e., by the annuity value of the increase in lifetime wealth). In the simple case in which the interest rate always equals the constant rate of time preference, the change in consumption exactly equals the change in permanent income. In the more general case in which r varies over time, the initial response of consumption may be greater or less than the change in permanent income, depending on the relation between the rates of interest and time preference in the current period and in all future periods. If we observe many income shocks, however, we can expect to see an approximate one-to-one relation on average. Nevertheless, Deaton (1987) reported that aggregate consumption changes by significantly less than permanent income in response to shocks to current income. This behavior has been dubbed the Deaton paradox, and it seems strong evidence against PILCH. In fact, the Deaton paradox is not evidence against PILCH or even paradoxical; rather, it is optimal behavior by the representative agent.

For ease of exposition, suppose the economy is in steady state when it experiences at time t a permanent proportional productivity shock. The restriction to permanent shocks is unimportant; the same results emerge for temporary shocks. Figure 1 shows the two possible responses. The response in Figure 1A occurs if the shock's wealth effect dominates the substitution effect; the response in Figure 1B occurs in the opposite case. In the following discussion, I restrict attention to the case shown in Figure 1A; the conclusions are even stronger for the case shown in Figure 1B. In Figure 1A, the optimal path for the economy is to jump down instantly to D_A and then move along the stable dynamic adjustment path to the new steady state, E_1. Consumption therefore jumps up initially and continues to rise thereafter until the steady state is reached.

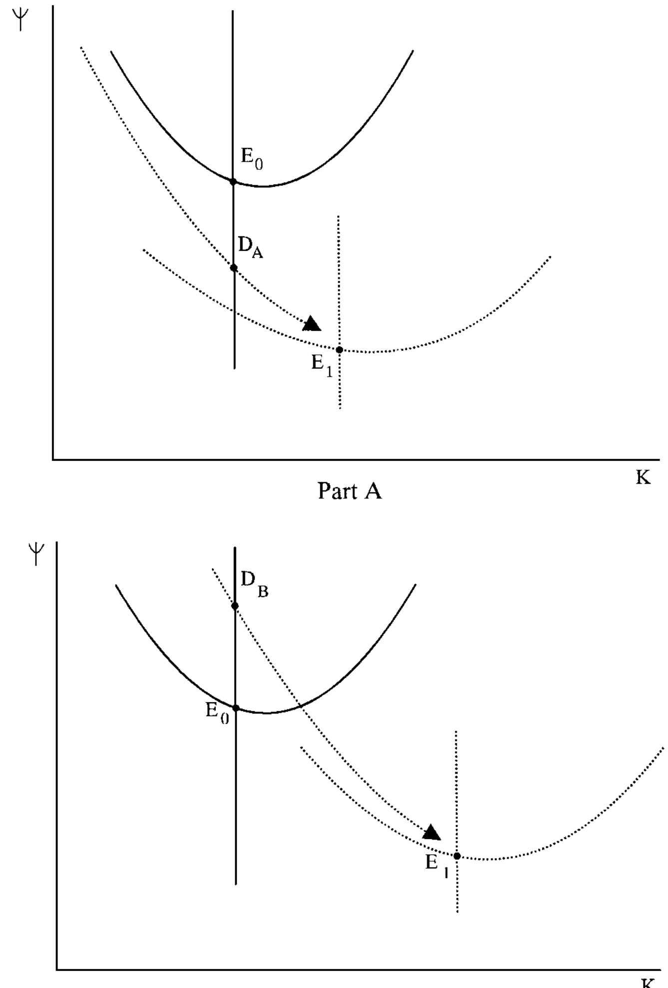

FIGURE 1. Dynamic adjustment of the representative agent to a productivity shock: (A) dominant wealth effect, (B) dominant substitution effect.

At the original capital stock k_t, current income y_t rises in response to the productivity shock, and it continues to rise as capital is accumulated. At time t, permanent income y_t^* rises by more than current income y_t because the future increases in current income induced by the shock are completely forecastable: $\Delta y_t < \Delta y_t^*$. At t, c_t rises by less than the increase in current income because there must be an increase in investment at t to make the capital stock subsequently grow as optimality requires: $\Delta c_t < \Delta y_t$. So, at time t, we observe $\Delta c_t < \Delta y_t < \Delta y_t^*$; that is, consumption rises by less than permanent income in response to the productivity shock. This behavior is precisely that of the Deaton paradox, but there is nothing paradoxical about it. It is merely optimal general-equilibrium behavior. Indeed, the economy's optimal general-equilibrium response is what makes a current income shock have positive aftershocks: Capital is accumulated and current income grows correspondingly. If movement toward the new steady state is fast enough, the time-series model estimated for income would be of low order, perhaps even ARIMA (0, 1, 1) as often found in the literature.[11] Part of the problem with the logic of the Deaton paradox is that it takes the behavior of aggregate income as exogenous to aggregate consumption when in fact the two are determined simultaneously as part of the representative agent's optimal response to an income shock. Because Deaton paradox behavior is predicted—indeed, required—by the theory, its actual occurrence in the data constitutes support for PILCH, not rejection.[12]

4.7. Liquidity Constraints

An atomistic individual facing a binding liquidity constraint cannot smooth his current income by borrowing and so consumes less than he desires. Anything that relaxes the constraint increases his consumption, and so, we expect the consumption by liquidity-constrained individuals to respond more strongly to changes in current income than does the consumption by unconstrained individuals. Averaging constrained and unconstrained individuals should produce a representative agent who is constrained and whose consumption therefore should be correlated with his current income, contrary to the implications of PILCH. The significant positive correlation between consumption and income in the aggregate data often has been interpreted as evidence of important liquidity constraints [e.g., Flavin (1985)].

There are two problems with this argument. First, as we already have seen, coefficient estimates obtained from the standard consumption function are econometrically unreliable, and so, finding a significant relation between current consumption and current income in the aggregate data proves nothing about the importance or existence of liquidity constraints. Second, the argument ignores the constraints of general equilibrium, which suggest that the quantitative impact of liquidity constraints will be nil in any case.

Suppose that there is no capital or investment so that all loans are consumption loans. Suppose also that the population is divided into two groups, B and L, which have identical preferences but different income paths so that, in the absence of liquidity constraints, group B borrows from group L. If we now impose liquidity constraints, group B's borrowing falls, and so its consumption also falls. At the same time, however, group L's lending falls, and its consumption must rise commensurately. Liquidity constraints change the distribution of aggregate consumption but cannot change its level because all current income is consumed in the aggregate (ignoring, as usual, second-order distribution effects). Consequently, an aggregate shock must have the same effect on the level of aggregate consumption whether liquidity constraints are present or not, and liquidity constraints therefore cannot introduce any new sensitivity of aggregate consumption to current income.

It is unclear whether these conclusions hold when capital and investment are introduced, but the simple example at least suggests the possibility that liquidity constraints do not have important effects on aggregate consumption in general equilibrium.[13]

4.8. Economic Growth and Age-Consumption Profiles

Finite (expected) lives and technical progress produce an upward trend in aggregate consumption. Young agents have higher permanent incomes than old agents because the young will enjoy the benefits of future economic growth whereas the old will not. We therefore have the time-series implication that the average agent's income and consumption both should grow as time passes. However, there also is an apparent cross-sectional implication. At any moment, the young should consume more than the old because the young have higher permanent incomes than the old. Consequently, we expect to see a downward-sloping relation between consumption and age, as shown in Figure 2A. A corollary is that, if we compare the age-consumption profiles for different countries, the profile should be more negatively sloped the higher a country's growth rate. Carroll and Summers (1991) test this implication and find no such relation. The countries they examine have very different growth rates but approximately the same hump-shaped age-consumption profiles. See Figure 2B. This cross-country age-consumption profile evidence seems to reject PILCH, but once again, that is merely an appearance arising from a fallacy of composition. PILCH does not predict that economic growth will cause a downward-sloping age-consumption profile, as a simple example shows.

Consider an economy with just two agents, one young and one old, who each live two periods. Suppose that they have the same production function and therefore the same current income at any given moment, that exogenous technical progress doubles productivity each period, and that technical progress is the only source of income variation. In the current period, both agents produce 100 units of output. Last period, when the current old agent was young, he produced 50 units of output. His permanent income was 75 (ignoring interest) when he was young. The currently young agent will produce 200 next period, and so, his permanent income is 150. Nevertheless, in the current period, each agent consumes his current income of 100 because young agents never can borrow against their expected future income. The current old agent will not lend to the young agent because next period the old agent will not be alive to be repaid; the current young agent cannot borrow from next period's young because they are not yet born. There is no way for any generation to borrow against its higher expected future income, and so, each generation consumes its current income each period. This is just a variant of the result first established by Samuelson (1958), and, like his, it generalizes to a model with generations living longer than two periods.

In the foregoing example, there is no borrowing or lending at all because the only variation in income is the intergenerational variation caused by economic growth. Everyone has the same current income (i.e., production function), and so,

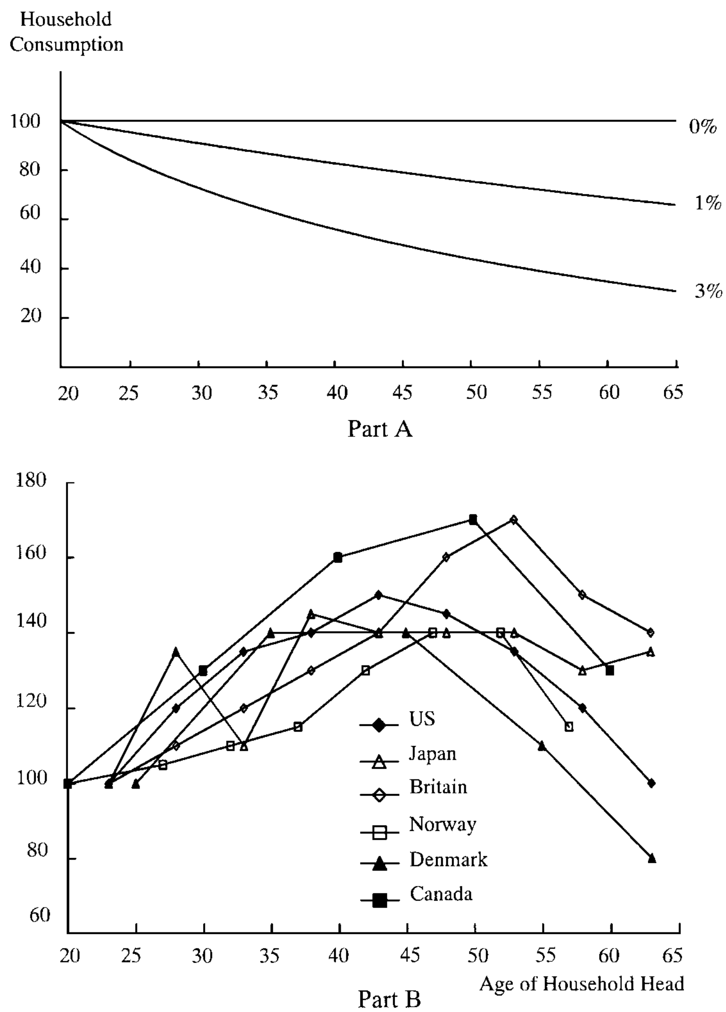

Figure 2. Theoretical (A) and empirical (B) age-consumption profiles. Consumption has been normalized to 100 for households with heads aged 20 [adapted from Carroll and Summers (1991)]. The three curves in Part A are the profiles corresponding to GDP annual growth rates of 0, 1, and 3 percent.

everyone also has the same consumption in any given period. In such a case, the economy's age-consumption profile always is flat. We can introduce the possibility of borrowing and lending by allowing other sources of income variation. If individuals in each generation have transitory income of different magnitudes occurring at different times, the usual incentives for borrowing and lending are present, and even cross-generational borrowing and lending will occur. That use of the capital market helps smooth consumption over time in accord with PILCH.

However, the variation in income arising from economic growth still cannot be smoothed by recourse to the capital market for the same reason as in the simpler example. Even though people have different current incomes, the economy's age-consumption profile will be flat as long as the income differences are randomly distributed across generations.

The important implication of these results is that countries differing only in their growth rates will have age-consumption profiles of approximately the same shape. Carroll and Summers's finding that countries indeed have similar-shaped profiles does not contradict PILCH but rather is consistent with it.[14]

4.9. Time-Series Behavior

One can derive predictions from PILCH about the time-series behavior of individual and aggregate consumption. Once again, the existing literature has not distinguished properly between atomistic and aggregate behavior and mistakenly has used the predictions for the former as a basis for analyzing the latter.

Consider the following version of PILCH, based on Friedman (1957):

$$c_t = c_{pt} + a_t,$$
$$c_{pt} = \beta y_{pt} \qquad 0 < \beta < 1,$$
$$y_t = y_{pt} + b_t,$$
$$E[a_t] = E[b_t] = E[y_{pt} b_t] = E[c_{pt} a_t] = E[a_t b_t] = 0,$$

where c_{pt} and y_{pt} are the individual's permanent consumption and income and a_t and b_t are his transitory consumption and income, respectively. Suppose permanent income follows the ARIMA (0, 1, 1) process

$$y_{pt} = y_{pt-1} + (1 + \alpha L) d_t$$

with $|\alpha| < 1$ and d_t white noise (as found in time-series studies of aggregate income), and suppose also that transitory consumption a_t is white noise. Falk and Lee (1990) show that consumption, apart from trend, must be the ARIMA (0, 1, 1) process

$$c_t - c_{t-1} = (1 - \gamma L) g_t, \tag{36}$$

where g_t is a random disturbance; α, γ, var (a_t), var (d_t), and var (g_t) are related as

$$\begin{aligned} \beta^2(1+\alpha^2)\sigma_d^2 + 2\sigma_a^2 &= (1+\gamma^2)\sigma_g^2, \\ \beta^2 \alpha \sigma_d^2 - \sigma_a^2 &= -\gamma \sigma_g^2; \end{aligned} \tag{37}$$

and the first-order autocovariance of Δc_t is given by

$$\operatorname{cov}(\Delta c_t, \Delta c_{t-1}) = \beta^2 \alpha \sigma_d^2 - \sigma_a^2.$$

Falk and Lee report empirical results that reject this ARIMA (0, 1, 1) representation for quarterly aggregate consumption data.

Although Falk and Lee's conclusions are correct for the individual, they do not apply to the representative agent. From (37), the coefficient γ in (36) depends on β. In both the atomistic and representative-agent problems, β depends on the interest rate r. For an atomistic individual, r and therefore β are exogenous, perhaps even constant. For the representative agent, however, r is not exogenous and generally not constant, and so, β is not constant, either. Consequently, γ is a time-varying parameter when (36) is applied to aggregate data, which means (36) is not truly ARIMA. If Δc is stationary for aggregate consumption data, the Wold theorem assures us it has an ARIMA representation, but (36) is not the right one. Falk and Lee's empirical rejection of (36) for aggregate data may reflect that fact.[15]

4.10. Open Economies

The foregoing results do not apply to a small open economy facing perfect markets. Such an economy can freely borrow and lend on the world capital market, thus breaking the tight link between its current income and its expenditure opportunities that has driven the analysis to this point. For such an economy, PILCH's implications are the same as for an atomistic agent, and so, we may be able to test PILCH by seeing if small open economies behave as atomistic agents. This possibility may be impractical. Real interest rates seem to differ persistently across countries, suggesting some kind of market imperfection such as information costs or differential tax and regulatory treatment of domestic vs. foreign investment. In the face of market imperfections, even small open economies may have the ability to affect their domestic interest rates by their choice of investment expenditure paths. In that case, the atomistic agent's behavior does not generalize to the aggregate economy, and we are back to our earlier analysis.

An interesting possibility is to turn the tables and use PILCH's implications to test for market perfection. A small open economy can exhibit aggregate consumption behavior consistent with that of an atomistic agent only if it operates in nearly perfect capital markets. If we test PILCH's atomistic-agent implications with such an economy's aggregate consumption data and cannot reject them, then we have evidence that the international capital market facing that country must be nearly perfect.

5. CONCLUSION

The unifying theme of this paper is that testing PILCH with aggregate consumption data imposes requirements and special considerations not relevant to tests based on microdata. Aggregate consumption data arise from the general-equilibrium behavior of the economy; equivalently, they represent the optimizing behavior of the representative agent. Representative agents do not behave as atomistic agents because representative agents must take into account the restrictions imposed by general equilibrium, which atomistic agents ignore. Aggregate and micro data thus have different characteristics, and PILCH has different implications for them. Most of the existing literature has ignored those differences and consequently has produced misspecified tests that are uninformative of the validity of PILCH.

Euler-equation tests of PILCH based on aggregate data suffer from possible inconsistency with GARP, possible instability of the aggregator function, and possible invalidity of the joint hypothesis of strong rational expectations needed to make the Euler-equation method operational. The traditional aggregate consumption-function tests are in some ways superior to Euler-equation tests because they require fewer and weaker identifying restrictions. However, most attempts to use such tests have tested the atomistic implications of PILCH with aggregate data and thus are biased, inconsistent, and inefficient. Famous rejections of PILCH, such as excess sensitivity of consumption to current income, are not valid rejections but rather are merely uninformative. Other supposed rejections, such as the Deaton paradox, actually support PILCH. There are valid aggregate consumption-function implications of PILCH, but those have not yet been tested.

The overall conclusion is that the aggregate and micro evidence on PILCH may not be as contradictory as they seem. Most of the aggregate evidence may simply be invalid.

NOTES

1. Derivations are in an Appendix available from the author on request.

2. Two possibilities are that (1) all agents have identical homothetic utility functions or (2) agents have homothetic utility functions that need not be identical, but also the relative income distribution is fixed and independent of prices. See Kirman (1992) and references cited therein.

3. Attanasio and Weber (1993), for example, argue that the failure of exact aggregation is a major reason for differences between Euler-equation parameter estimates obtained from micro- and macrodata.

4. Even buying economic analysis from experts seems unlikely to be informative for most people, for the experts often disagree wildly in their predictions.

5. Indeed, Goodfriend (1992) and Deaton (1992) have shown formally that lags in collecting aggregate information introduces correlations between current consumption and lagged values of income.

6. My own personal experience of several years in the Federal Reserve system, where I observed policy making by all branches of government, is that policy choices change often but policy procedures change rarely, if at all.

7. A partial exception arises from uncertainty. How consumers respond to uncertainty depends on whether their marginal utility is convex or concave in consumption. There is evidence of convex marginal utility, which implies cautious consumers (Caballero, 1990), and one may want to impose convex marginal utility on the utility function. Doing so is still much less restrictive than imposing a specific functional form.

8. See the Appendix available from the author for the derivation of this and all other results discussed in the text.

9. “Rule-of-thumb” consumers are those who consume their current income at all times. The name is from Campbell and Mankiw (e.g., 1990); the concept goes back at least to Mayer (1972), who called it the standard income hypothesis.

10. In fact, the results are even stronger. It can be shown that PILCH places no restrictions on either the sign or significance of any of the coefficients in (35), even though it does restrict the coefficients in (34).

11. Rossana and Seater (1995) show that the low order obtained for aggregate annual income is an artifact of temporal aggregation; the true order is quite a bit higher. Consequently, the approach to the steady state need not be very fast to be consistent with the data.

12. These results were foreshadowed by Christiano (1987), who obtained Deaton-paradox behavior from interest-rate movements in simulations of a real business-cycle model. The results herein are much more general than Christiano's, for they are analytical and apply to any competitive macroeconomy rather than being the outcome of a particular simulation model.

13. See Scheinkman and Weiss (1986) for an analysis of liquidity constraints in general equilibrium. They do not distinguish between distributional and other effects and do not address explicitly the effect of liquidity constraints on aggregate consumption's sensitivity to aggregate shocks, and so, there is room for further analysis of liquidity constraints' general-equilibrium effects.

14. The humped shape of the age-consumption profiles does disagree with the limited version of PILCH discussed here, but apparently extensions that account for uncertainty (Caballero, 1990) and changing family size (Attanasio and Browning, 1995; Attanasio and Weber, 1995) can account for it.

15. A similar problem applies to the implications of HARA utility for consumption. When utility is HARA, an individual's consumption is linearly related to his wealth. This implication does not carry over to the aggregate data. The result for the individual depends on the interest rate being independent of the individual's wealth. In contrast, the interest rate does depend on the representative individual's wealth (the capital stock), and the linear relation between consumption and wealth does not hold because of the extra effect that changes in wealth have on the interest rate.

REFERENCES

Attanasio, O.P. & M. Browning (1995) Consumption over the life cycle and over the business cycle. *American Economic Review* 85, 1118–1137.

Attanasio, O.P. & G. Weber (1993) Consumption growth, the interest rate and aggregation. *Review of Economic Studies* 60, 631–649.

Attanasio, O.P. & G. Weber (1995) Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey. *Journal of Political Economy* 103, 1121–1157.

Caballero, R.J. (1990) Consumption puzzles and precautionary savings. *Journal of Monetary Economics* 25, 113–136.

Campbell, J.Y. & N.G. Mankiw (1990) Permanent income, current income, and consumption. *Journal of Business and Economic Statistics* 8, 265–279.

Carroll, C.D. & L.H. Summers (1991) Consumption growth parallels income growth: Some new evidence. In B.D. Bernheim & J.B. Shoven (eds.), *National Saving and Economic Performance*, pp. 305–307. Chicago: University of Chicago Press.

Christiano, L.J. (1987) Why is consumption less volatile than income? *Quarterly Review, Federal Reserve Bank of Minneapolis Quarterly Review* 11, (Fall) 1987, 2–20.

Cooley, T.F., S.F. LeRoy, & N. Raymon (1984) Econometric policy evaluation: Note. *American Economic Review* 74, 467–470.

Deaton, A. (1987) Life-cycle models of consumption: Is the evidence consistent with the theory? In T.F. Bewley (ed.), *Advances in Econometrics—Fifth World Congress*, Vol. II, pp. 121–148. Cambridge, England: Cambridge University Press.

Deaton, A. (1992) *Understanding Consumption*. Oxford: Oxford University Press.

Eichenbaum, M. & L.P. Hansen (1990) Estimating models with intertemporal substitution using aggregate time series data. *Journal of Business and Economic Statistics* 8, 53–69.

Falk, B. & B.-S. Lee (1990) Time-series implications of Friedman's permanent income hypothesis. *Journal of Monetary Economics* 26, 267–283.

Feige, E.L. & D.K. Pearce (1976) Economically rational expectations: Are innovations in the rate of inflation independent of innovations in measures of monetary and fiscal policy? *Journal of Political Economy* 84, 499–522.

Flavin, M.A. (1985) Excess sensitivity of consumption to current income: Liquidity constraints or myopia? *Canadian Journal of Economics* 18, 115–136.

Fleissig, A.R., A. Hall & J.J. Seater (1997) GARP, separability, and the representative agent. Typescript.

Fleissig, A.R., A.R. Gallant & J.J. Seater (1997) Theory matters: GARP, separability, aggregation, and Euler equation estimation. Typescript.

Friedman, M. (1957) *A Theory of the Consumption Function*. Princeton: Princeton University Press.

Gallant, A.R. & G. Tauchen (1989) Semi-nonparametric estimation of conditionally constrained heterogeneous processes: Asset pricing implications. *Econometrica* 57, 1091–1120.

Geweke, J. (1985) Macroeconometric modeling and the theory of the representative agent. *American Economic Review* 75, 206–210.

Goodfriend, M. (1992) Information-aggregation bias. *American Economic Review* 82, 508–519.

Hall, R.E. (1978) Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy* 86, 971–987.

Hansen, L.P. & K.J. Singleton (1983) Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy* 91, 249–265.

Jeong, J. & G.S. Maddala (1991) Measurement errors and tests for rationality. *Journal of Business and Economic Statistics* 9, 431–439.

Kirman, A.P. (1992) Whom or what does the representative agent represent? *Journal of Economic Perspectives* 6, 117–136.

Lovell, M. (1986) Tests of the rational expectations hypothesis. *American Economic Review* 76, 110–124.

Lucas, R.E., Jr. (1976) Econometric policy evaluation: A critique. In K. Brunner & A.H. Meltzer (eds.), *The Phillips Curve and Labor Markets*, pp. 19–46. New York: North-Holland.

Machina, M.J. (1987) Choice under uncertainty: Problems solved and unsolved. *Journal of Economic Perspectives* 1, 121–154.

Mayer, T. (1972) *Permanent Income, Wealth, and Consumption: A Critique of the Permanent Income Theory, the Life-Cycle Hypothesis, and Related Theories*. Berkeley: University of California Press.

Meinhold, R.J. & N.D. Singpurwalla (1983) Understanding the Kalman filter. *American Statistician* 37, 123–127.

Pischke, J.-S. (1995) Individual income, incomplete information, and aggregate consumption. *Econometrica* 63, 805–840.

Rossana, R.J. & J.J. Seater (1995) Temporal aggregation and economic time series. *Journal of Business and Economic Statistics* 13, 441–451.

Salop, S.C. (1987) Evaluating uncertain evidence with Sir Thomas Bayes: A note for teachers. *Journal of Economic Perspectives* 1, 155–160.

Samuelson, P.A. (1958) An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of Political Economy* 66, 467–482.

Scheinkman, J.A. & L. Weiss (1986) Borrowing constraints and aggregate economic activity. *Econometrica* 54, 23–45.

Varian, H. (1982) The nonparametric approach to demand analysis. *Econometrica* 50, 945–973.

Varian, H. (1983) Nonparametric tests of consumer behavior. *Review of Economic Studies* 50, 99–110.

Zarnowitz, V. (1992) *Business Cycles: Theory, History, Indicators, and Forecasting*. Chicago: University of Chicago Press for NBER.