

Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)

Christopher D. Carroll
ccarroll@jhu.edu

The Johns Hopkins University

September 7, 1997

Abstract

This paper shows that standard empirical methods of estimating log-linearized consumption Euler equations cannot successfully uncover structural parameters like the coefficient of relative risk aversion from a dataset of simulated consumers behaving exactly according to the standard model. Furthermore, consumption growth for the simulated consumers is very highly statistically related to predictable income growth – and thus standard ‘excess sensitivity’ tests would reject the hypothesis that consumers are behaving according to the standard model. Results are not much better for the second-order approximation to the Euler equation. The paper concludes that empirical estimation of consumption Euler equations should be abandoned, and discusses some alternative empirical tests of consumption models that are not subject to the problems of Euler equation estimation.

Keywords: Euler equation, uncertainty, consumption, excess sensitivity

JEL Classification Codes: C6, D91, E21

Thanks to participants in the 1997 NBER Summer Institute Consumption workshop, and to Orazio Attanasio, Martin Browning, John Campbell, Sydney Ludvigson, and Chris Paxson.

1 Introduction

Estimation of Euler equations has dominated empirical research on consumption over the almost twenty years since Hall's (1978) seminal work. Unfortunately, despite scores of empirical studies using household data, Euler equation estimation has not fulfilled its early promise to reliably uncover preference parameters such as the intertemporal elasticity of substitution. Even more frustrating, the model does not even seem to fail in a consistent way: some studies find strong evidence of 'excess sensitivity' of consumption to predictable income growth, while others find much less, or even no, excess sensitivity.

This paper offers an explanation for the conflicting and inconsistent empirical results, by showing that when the Euler equation estimation methods most commonly used on household data are applied to simulation data generated by a standard model of consumption under uncertainty, those methods are incapable of producing a consistent estimate of the intertemporal elasticity of substitution. Furthermore, 'excess sensitivity' tests can find either high or low degrees of sensitivity, depending on the exact nature of the test.

All of the theoretical problems stem from approximation error. The standard procedure has been to estimate a log-linearized version of the Euler equation. The paper shows, however, that the higher-order terms are endogenous with respect to the first-order terms (and also with respect to omitted variables), rendering consistent estimation of the log-linearized Euler equation impossible. Unfortunately, the second-order approximation fares only slightly better. The paper concludes that empirical estimation of approximated consumption Euler equations should be abandoned, and discusses some alternative empirical tests of consumption behavior that are not subject to the problems of Euler equation estimation.

The paper begins by presenting the specific version of the dynamic optimization problem that is tested. The next section describes the standard empirical methodology for estimating Euler equations and summarizes the results that have been reported in the literature. The next section describes the details of the simulations which generate the data which are to be analyzed using the same methods used in the empirical literature.

The following section is the heart of the paper: it shows that the standard empirical methods cannot produce consistent estimates of true model parameter values. The next section describes several empirical strategies that are candidates to replace Euler equation estimation, and the final section concludes.

2 The Model

Assume that the consumer is solving the following maximization problem (essentially the same as the model in Carroll (1992, 1997) and Zeldes (1989)):

$$\begin{aligned}
\max_{C_t} \quad & u(C_t) + E_t \sum_{s=t+1}^T \beta^{(s-t)} u(C_s) \\
s.t. \quad & X_{t+1} = R(X_t - C_t) + \tilde{Y}_{t+1} \\
& \tilde{Y}_s = \tilde{P}_s \tilde{V}_s \\
& \tilde{P}_s = G P_{s-1} \tilde{N}_s \\
& u(C) = \frac{C^{(1-\rho)}}{1-\rho} \text{ where } \rho > 1
\end{aligned} \tag{1}$$

where P_s is permanent labor income, which is buffeted by lognormally distributed mean-one shocks \tilde{N}_s with variance of $\log N = \sigma_n^2$, implying that $\log P_s$ follows a random walk with drift; Y_s is current labor income, which is equal to permanent labor income multiplied by a mean-one transitory shock \tilde{V} which is equal to zero with probability p and otherwise is distributed lognormally with variance of $\log V = \sigma_v^2$, and with a mean that guarantees that $E_t \tilde{V}_{t+1} = 1$; the interest rate, the growth rate of income, and the time preference factor, respectively R , G , and β , are constant; and the consumer's utility function is of the Constant Relative Risk Aversion form with coefficient of relative risk aversion $\rho \geq 1$.

The solution to this model obeys the Euler equation:

$$R\beta E_t [C_{t+1}/C_t]^{-\rho} = 1. \tag{2}$$

As written, this problem has two state variables, the level of liquid assets and the level of permanent labor income. Carroll (1996) shows that this problem can be converted to a single-state-variable problem by dividing through by the level of permanent income P_t ,

implying that at each age of life there is an optimal rule relating the ratio of cash-on-hand to permanent income $x_t = X_t/P_t$ to the ratio of consumption to permanent income $c_t = C_t/P_t$.

The model is solved numerically by backwards induction on the Euler equation. In the last period of life, the optimal plan is to consume everything, $c_T^*(x_T) = x_T$. In the next-to-last period, designating $t = T - 1$, the standard Euler equation for marginal utility is

$$R\beta E_t \left[\frac{P_{t+1}c_{t+1}^*(x_{t+1})}{P_t c_t} \right]^{-\rho} = 1. \quad (3)$$

For a given value of x_{T-1} this equation can be solved numerically to find the optimal value of c_{T-1} . This is done for a grid of possible values for x_{T-1} and a numerical optimal consumption rule $c_{T-1}^*(x_{T-1})$ is constructed by linear interpolation between these points. Given $c_{T-1}^*(x_{T-1})$ the same methods can be used to construct $c_{T-2}^*(x_{T-2})$ and so on to any arbitrary number of periods from the end of life.¹ Carroll (1996) shows that if Deaton's 'impatience' condition $R\beta E_t(G\tilde{N}_{t+1})^{-\rho} < 1$ holds, these successive optimal consumption rules will converge to a steady-state rule as the horizon recedes, and that consumers behaving according to the converged rule can be described as engaging in 'buffer-stock' saving. I will denote the optimal consumption rule for any period t as $c_t^*(x_t)$ and the converged consumption rule as $c^*(x)$.

To verify the accuracy of the numerical solution, Figure 1 plots $R\beta E_t \left[\frac{P_{t+1}c_{t+1}^*(x_{t+1})}{P_t c_t^*(x_t)} \right]^{-\rho}$ as a function of x_t . Approximation errors will lead the function to differ from one at points away from the gridpoints chosen for x_t . The figure shows that the errors involved in numerical solution of the model are very small; the function is so close to one over the entire plotted range (which encompasses the range of values of wealth that actually arise when the model is simulated) that it appears to be a solid line exactly at one.

¹For more details on the method of solution, see Carroll (1992, 1997).

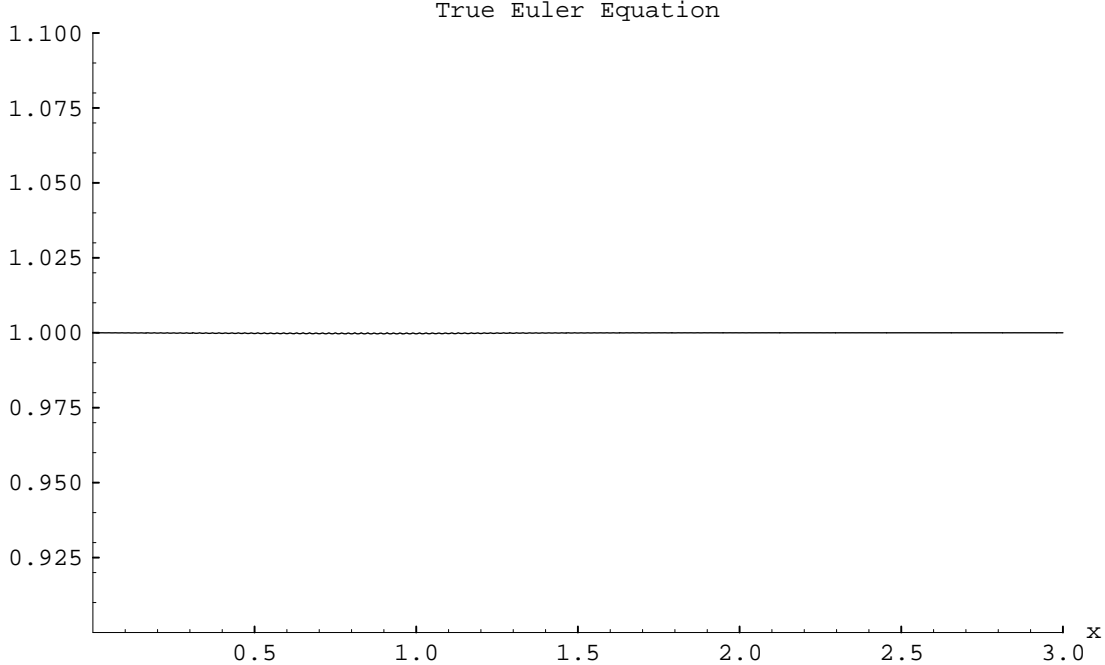


Figure 1: Numerical Value of $R\beta E_t \left[\frac{P_{t+1}c^*(x_{t+1})}{P_t c^*(x_t)} \right]^{-\rho}$

3 The Standard Procedure

3.1 Derivation of the Log-Linearized Consumption Euler Equation

The “Log-Linearized” consumption Euler equation of this paper’s title is obtained by taking a first-order Taylor expansion of the nonlinear Euler equation (2), and making some approximations. For every possible C_t and C_{t+1} there will be some η_{t+1} for which $C_{t+1} = (1+\eta_{t+1})C_t$ (assuming that consumption is always positive). Since we rarely expect to see consumption rise or fall dramatically from period to period, it seems reasonable to use the approximation $(1+\eta_{t+1})^{-\rho} \approx 1 - \rho\eta_{t+1}$ which corresponds to the first-order Taylor expansion of $(1+\eta_{t+1})^{-\rho}$ around the point $\eta_{t+1} = 0$. The Euler equation (2) then becomes:

$$R\beta E_t(1 - \rho\eta_{t+1}) \approx 1. \tag{4}$$

A simple transformation of this first-order approximation to the Euler equation has been the basis for most of the estimation of consumption Euler equations. By definition $1 + \eta_{t+1} = C_{t+1}/C_t$ and using the approximation that for ‘small’ ϵ , $\log(1 + \epsilon) \approx \epsilon$ we obtain

$\eta_{t+1} \approx \log C_{t+1} - \log C_t = \Delta \log C_{t+1}$. Substituting this back into equation (4) gives

$$R\beta(1 - \rho E_t \Delta \log C_{t+1}) = 1. \quad (5)$$

Finally, taking the log of both sides, implicitly defining the time preference rate δ from $\beta = 1/(1+\delta)$ so that $\log R\beta \approx r - \delta$, and using the approximation $\log(1 - \rho E_t \Delta \log C_{t+1}) \approx -\rho E_t \Delta \log C_{t+1}$ gives

$$\begin{aligned} (r - \delta) - \rho E_t \Delta \log C_{t+1} &\approx 1 \\ E_t \Delta \log C_{t+1} &\approx \rho^{-1}(r - \delta), \end{aligned} \quad (6)$$

or, defining the expectation error $\epsilon_{t+1} = \Delta \log C_{t+1} - E_t \Delta \log C_{t+1}$, an alternative way to express this result is:

$$\Delta \log C_{t+1} \approx \rho^{-1}(r - \delta) + \epsilon_{t+1} \quad (7)$$

where ϵ_{t+1} is iid and the law of iterated expectations implies that it is uncorrelated with any variable known at time t (Hall (1978)).

Those authors made uncomfortable by the first-order approximations involved in deriving equation (7) have sometimes been reassured by a well-known result that suggests that the second-order approximation leads to the same estimating equation. The second-order Taylor approximation $(1 + \eta_{t+1})^{-\rho}$ around $\eta_{t+1} = 0$ is $(1 + \eta_{t+1})^{-\rho} \approx 1 - \rho\eta_{t+1} + \frac{\rho(\rho+1)}{2}\eta_{t+1}^2$. Solving for $\Delta \log C_{t+1}$ as above, the end result is

$$\Delta \log C_{t+1} \approx \rho^{-1}(r - \delta) + \frac{\rho + 1}{2} E_t \eta_{t+1}^2 + \epsilon_{t+1}, \quad (8)$$

and if η_{t+1}^2 is uncorrelated with r and δ , then the estimating equation (8) is still a valid way of estimating the value of ρ because the $E_t \eta_{t+1}^2$ term should be absorbed in the constant term of the regression.²

²A common alternative way of deriving essentially the same result is to assume that the consumption shocks are lognormally distributed and independent of the other variables in the model; in that case the last term in equation (8) is the variance of the consumption innovations rather than the square, and its coefficient is $\rho/2$ rather than $(\rho + 1)/2$.

3.2 Previous Empirical Results

To keep the notation simple, the derivations thus far have implicitly assumed that ρ , δ , and r are constants. Of course, if these parameters were constant across all times, places, and people then it would be impossible to estimate a coefficient ρ in an equation like (7). In practice, Euler equations like (7) have mainly been estimated in two ways. In microeconomic data, the most common procedure has been to estimate the equation across different consumers at a point in time, by identifying groups of consumers for whom different interest rates apply. In macroeconomic data, the equation has been estimated by exploiting time-variation in the aggregate interest rate.³ The principal purpose of this paper is to show that the usual procedures for microeconomic estimation of this equation do not work; it seems likely that similar problems will apply to estimation based on time-series variation in interest rates, but I leave that question for future research.

The instrumental variables approach to estimating the model using microeconomic data can be thought of as equivalent to taking means across groups of consumers with different characteristics. For example, typical instruments used in the empirical literature are education group or occupation. Denoting distinct groups identified by the instruments as j , and denoting the mean consumption growth for members of group j as $(\Delta \log C_{t+1})_j$ and the group-specific values of the parameters as ρ_j , r_j , and δ_j ,⁴ equation (7) becomes:

$$(\Delta \log C_{t+1})_j \approx \rho_j^{-1}(r_j - \delta_j) + (\epsilon_{t+1})_j \quad (9)$$

Thus, the standard log-linearized empirical Euler equation boils down to:

$$(\Delta \log C_{t+1})_j = \alpha_0 + \alpha_1 r_j + (\epsilon_{t+1})_j \quad (10)$$

where the understanding has been that α_1 , the coefficient on r , should be a consistent estimate of the intertemporal elasticity of substitution, ρ^{-1} . This will be true if three conditions hold: first, the approximations involved in deriving equation (7) are not problematic; second, any differences in δ_j across groups are uncorrelated with whatever differences

³A few studies have had enough cross-sections of household data to exploit time-variation in the aggregate interest rate using household data. See in particular Attanasio and Weber (1985).

⁴Assume for convenience that everyone in group j has the same values for these parameters.

there may be in r_j ; and, finally, there are no differences across groups in ρ_j .

Empirical results for estimating equations like (10) have been poor. Usually the α_1 term is estimated to be insignificantly different from zero; only a few studies have found significantly positive values of ρ .⁵ However, the poor results in estimating ρ often been interpreted as reflecting poor identifying information about differences in r across groups, rather than as important rejections of the Euler equation itself.⁶

The potential empirical problems with identifying exogenous variation in interest rates across households have led many papers to focus on another feature of the model: Hall’s ‘random walk’ proposition. Hall showed that in a model with quadratic utility, consumption should follow a random walk and no information known at time t should help to forecast consumption growth between t and $t + 1$. The alternative hypothesis has usually been that consumption is ‘excessively sensitive’ to forecastable income growth. Formally, denoting the average growth rate of income for consumers in group j as $(\Delta \log Y_{t+1})_j$, the equation most commonly estimated has been:

$$(\Delta \log C_{t+1})_j = \alpha_0 + \alpha_1 r_j + \alpha_2 (E_t \Delta \log Y_{t+1})_j + \epsilon_j, \quad (11)$$

and the ‘random walk’ proposition implies that $\alpha_2 = 0$ when the expected growth rate of income is instrumented using information known by the consumers at time t .

Empirical results estimating equation (11) using micro data have been hardly better than those estimating the baseline equation (10).⁷ In a comprehensive survey article, Browning and Lusardi (1996) cite roughly twenty studies that have estimated the coefficient on predictable income growth. Estimates of the marginal propensity to consume out of predictable income growth ranged from zero (consistent with the CEQ LC/PIH model) up to 2. An optimist might note that most estimates are in the range between 0 and .6.

⁵See the survey paper by Browning and Lusardi (1996) for more details.

⁶Usually identification has been obtained by calculating a marginal tax rate for each person and using the variation in marginal tax rates across households to identify an after-tax interest rate. This is problematic if the level of income is correlated with tastes. One simple mechanism for such a correlation is capital accumulation: if patient consumers save more they will eventually have a higher level of capital income, generating a correlation between tastes and the marginal tax rate.

⁷Although, interestingly, when the equation is estimated using macro data it reliably generates a coefficient of around 0.5. See below for a potential explanation.

3.3 The Explanation?

Carroll (1992, 1996, 1997) has challenged the foregoing empirical methodology on the grounds that theory implies that the higher-order terms in the approximation cannot be ignored because they are *endogenous* and in particular are correlated with ρ_j, δ_j , and, most fatally, r_j and $(E_t \Delta \log Y_{t+1})_j$. Carroll (1992, 1996, 1997) shows that ‘impatient’ consumers behaving according to the standard CRRA intertemporal optimization model described above will engage in “buffer-stock” or target saving behavior.⁸ Those papers show that, among a collection of buffer-stock consumers with the same parameter values, if the distribution of x across consumers has converged to an ergodic distribution, then average consumption growth across the members of the group will be equal to average permanent income growth. Thus, if we have j groups of consumers such that within each group j all consumers have the same parameter values, and x has converged to its ergodic distribution within each group j , then

$$(\Delta \log C_{t+1})_j = (\Delta \log P_{t+1})_j = g_j. \quad (12)$$

The intuition for this result is fairly simple: If consumers are behaving according to a target-saving or buffer-stock model, then it is impossible for consumption growth to be permanently different from underlying income growth. If consumption growth were forever greater than income growth, consumption would eventually exceed income by an arbitrarily large amount, driving wealth to negative infinity. If consumption growth were permanently less than income growth, income would eventually exceed consumption by an arbitrarily large amount, driving wealth to infinity. Thus, in a model where there is an ergodic distribution of wealth across consumers, it is impossible for average consumption growth to differ permanently from average income growth.⁹

As an aid to understanding the nature of the endogeneity problem, suppose that the second-order approximation equation (8) captures all of the important endogeneity so that

⁸The term ‘impatient’ here and henceforth refers to the condition $R\beta E_t(G\tilde{N}_{t+1})^{-\rho} < 1$. Note that, so long as income is growing over time $G > 1$, consumers can be impatient in the required sense even if $\beta = 1$ so that they do not discount future utility at all.

⁹For much more careful discussion and arguments, see Carroll (1996, 1997).

the terms of third order and higher can safely be ignored (we will examine this assumption carefully below). Assume that ρ does not differ across the groups, and rewrite the second order approximation equation (8) in the new notation:

$$(\Delta \log C_{t+1})_j \approx \rho^{-1}(r_j - \delta_j) + \frac{1+\rho}{2}(E_t \eta_{t+1}^2)_j. \quad (13)$$

If the members of group j are distributed according to their ergodic distribution, it should be the case that the average value of η_{t+1}^2 is equal to the average of its expected value. Substituting $(\eta_{t+1}^2)_j$ for $(E_t \eta_{t+1}^2)_j$ in equation (13) we now have two equations, (13) and (12), for average consumption growth for members of group j . The only way both equations can hold simultaneously is if the $(\eta_{t+1}^2)_j$ term is an endogenous equilibrating variable; in particular, the two equations can be solved for the value this term must take:

$$(\eta_{t+1}^2)_j \approx \frac{2}{1+\rho}[g_j - \rho^{-1}(r_j - \delta_j)] \quad (14)$$

This equation makes abundantly clear the econometric problem with estimating the log-linearized Euler equation (7): $(\eta_{t+1}^2)_j$ is an omitted variable in the regression equation and *theory implies that it is correlated with r_j* (as well as with g_j , δ_j and ρ_j if they differ across groups). Hence it will be impossible to get a consistent estimate of the coefficient on r_j if the $(\eta_{t+1}^2)_j$ term is omitted from the equation.

The easiest way to understand how the mechanism works is to think of η_{t+1}^2 as a measure of the degree of undesirable variation in consumption growth caused by the uncertainty of income. Because consumers with less wealth have less ability to buffer consumption against income shocks,¹⁰ there will be a direct relationship between the level of wealth and the value of $E_t \eta_{t+1}^2$. In fact, the size of the target buffer stock of wealth is the real equilibrating factor in the model. For example, consumers who are more impatient (higher δ) will have a smaller value of the $\rho^{-1}(r_j - \delta_j)$ term in the Euler equation. However, impatient consumers will also hold less wealth, leading to a higher value of $E_t \eta_{t+1}^2$. Across steady-states, the higher value of the η_{t+1}^2 term should exactly offset the lower

¹⁰This is an implication of the concavity of the consumption function proven by Carroll and Kimball (1996).

value of the $\rho^{-1}(r_j - \delta_j)$ term, leaving the growth rate of consumption at g_j regardless of the value of δ_j (so long as the impatience condition is satisfied).¹¹

Another thought experiment illustrates the econometric problem very clearly. Consider a dataset composed of consumers who satisfy the impatience condition and thus are buffer-stock savers. Suppose these consumers are identical in every respect (including having a common expected growth rate of permanent income g) except that different consumers face different interest rates. Suppose further that the econometrician can observe each household's interest rate. If equation (10) were a valid empirical specification this would be the ideal dataset for estimating the intertemporal elasticity of substitution. But what happens when equation (10) is estimated on this dataset? The regression will estimate $\alpha_0 = g$ and $\alpha_1 = 0$ regardless of the true value of ρ , because the expected growth rate of consumption will be equal to g for all the consumers despite the difference in interest rates across groups. The reason is that the consumers facing a higher interest rate will hold more wealth, and therefore will have a lower value of $E_t \eta_{t+1}^2$ by an amount that exactly offsets the higher interest rate they face.

The foregoing theoretical arguments are not, in themselves, sufficient to definitively discredit the estimation of log-linearized Euler equations, because the arguments were predicated on two untested assumptions: that consumers within each group are distributed according to an ergodic distribution, and that the second-order approximation is not problematic. Only simulation methods can determine whether the behavior of the second-order approximation under the ergodicity assumption is a good or bad guide to the behavior of a finite collection of consumers obeying the model over limited time periods. I therefore turn now to simulations.

4 The Simulations

The procedure for generating simulated data from the model is as follows. First, I solve the model as specified above for the baseline set of parameter values indicated in Table 1,

¹¹This statement assumes that the second-order approximation holds exactly. The more general statement would be that all of the higher-order terms together should take on values that make $(\Delta \log C_{t+1})_j = g_j$.

Parameter	Low	Baseline	High
r	0.00	0.02	0.04
δ	0.00	0.04	0.08
g	0.01	0.03	0.05
ρ	1	3	5
σ_n	0.05	0.10	0.15
σ_v	0.05	0.10	0.15

Table 1: Parameter Values

yielding a baseline consumption rule $c^*(x)$. I then solve the model for two alternative values of each of the model’s parameters, leaving the other parameters fixed at their baseline levels. For example, I solve the model in the case where all parameter values are at their baseline levels except that the interest rate is assumed to be 0 percent, then I solve for the case where the interest rate is 4 percent. This generates two alternative consumption rules $c_{r=.00}^*(x)$ and $c_{r=.04}^*(x)$ where the subscripts indicate which parameter is being set to a value different from baseline.

When all of the optimal consumption rules have been generated, I perform the simulations. For each combination of parameter values (‘group,’ for short), I set up a population of one thousand consumers who begin ‘life’ with zero assets. For their first year of life, I draw random income shocks from the income distribution functions described above. I next use the appropriate consumption rule to determine first period consumption. First period’s income and consumption determine the savings with which the the consumers enter the second period; I draw random income shocks again, and again apply the consumption rule, yielding period two consumption and saving. I repeat this exercise for twenty periods (‘years’) in a row, discarding the first 9 periods in order to allow the distribution of x across consumers to ‘settle down’ to something approximating the ergodic distribution. For the baseline set of parameter values, Figure 2 plots the theoretical distribution of x after ten years of simulation against the ergodic distribution; the match is very close, suggesting that nine years of presample simulation are adequate preparation.

The data from years 10-20 are processed to generate 10,000 observations of $\Delta \log C_{t+1}$,




Figure 2: Distribution of x After 10 Years (Solid) vs. Ergodic Distribution (Dashing)

r , $\Delta \log Y_{t+1}$, and the dummy variables indicating group membership. With the exception of the interest rate, the data do not contain the actual values of the parameters; instead, they contain dummy variables for each parameter that equal one or zero for each consumer. Roughly speaking, these dummy variables correspond to the ‘instruments’ such as occupation, education, and race used in actual data.

The goal is to characterize the kinds of regression results that an econometrician would obtain using a sample of data drawn from these distributions. The appropriate strategy is therefore a Monte Carlo procedure which reports both the mean parameter estimates that would be obtained by a large number of studies on such data, and the variation in parameter estimates that would be found across the different studies.

My Monte Carlo procedure is as follows. For each ‘group’ to be included in a regression, I draw a random sample of 1000 observations from the 10,000 available for that group. I then perform the regressions and record the coefficient estimates and standard errors. I then draw another sample of 1000 observations for each group, perform another regression, and record the results. I repeat this procedure 100 times to obtain a distribution of

parameter estimates and standard error estimates.¹²

Note that there are several respects in which the ‘econometrician’ examining the simulated data is better off than his counterpart using actual data. First, there is no measurement error in the simulated data for either income or consumption; estimates of the fraction of measurement error in the PSID data on food consumption range up to 92 percent. Second, the econometrician working with simulated data can directly observe the interest rate that applies for each household. In empirical work there is rarely a really convincing way to identify exogenous differences in interest rates across the different households in the sample. Third, the different ‘groups’ in the simulations differ from the baseline parameter values in only a single dimension (parameter) at a time. In reality, occupation or education may be correlated with several parameters; for example, education is highly correlated with the growth rate of income, but may also be correlated with the time preference rate. Finally, the typical empirical dataset probably has fewer than a hundred consumers in any given instrumented age/occupation or age/education cell, while I have a thousand consumers for each possible combination of parameter values. The purpose of these simulations is to show that even in ideal circumstances, Euler equation estimation by standard IV methods does not work. Presumably there is even less reason to expect it to work under the less than ideal circumstances faced in actual data.

5 Estimating Consumption Euler Equations on the Simulated Data

5.1 The Log-Linearized Euler Equation

Table 2 presents the results when the log-linearized Euler equation (10) is estimated on the simulated data. The first row presents results when equal numbers of consumers from each

¹²Ideally, it would be better to estimate thousands of regressions on thousands of totally independently simulated consumers. However, creating 10,000 simulated consumers for each combination of parameter values is close to the limit of what is possible with the computer facilities at my disposal. The results I report are very similar to those obtained from several different runs of the complete set of programs, so it seems very unlikely that results would be much different if I were able to simulate more consumers.

possible parametric combination (except for deviations of ρ from baseline) are included;¹³ the second row presents results when only the consumers with baseline parameter values (group BASE) and those for whom the interest rate differs from the baseline (group R) are included. The second column indicates the set of instruments used for predicting all instrumented variables in the regression. Since r is the only explanatory variable included in the regressions reported in rows 1 and 2, the dummy variable indicating interest rate group (RDUM) is the only instrument that makes sense in these two regressions.

I exclude from the regressions all consumers for whom income was zero in either period of observation, $V_t = 0$ or $V_{t+1} = 0$, for two reasons. First, such data are typically excluded from the empirical regressions whose methods I am trying to duplicate. Second, extreme income shocks tend to interact strongly with the nonlinearities of the model, so even a relatively small number of such extreme events could heavily influence the results. It is therefore a more compelling indictment of the estimation method if it performs badly even when such extreme events are excluded.

As noted above, I estimate the regressions 100 times with 100 different collections of simulated consumers. For each variable, the table presents the mean (across the 100 regressions) of the coefficient estimates and of the standard errors. Next to the means are the fifth and ninety-fifth percentiles in the distribution of coefficient estimates and standard error estimates. The last column indicates the average number of observations in each regression. Because the probability that either $V_t = 0$ or $V_{t+1} = 0$ is .01, this number should on average be equal to $0.99 \cdot 1000 \cdot (\text{number of groups included in regression})$. For example, one would expect a sample size of $0.99 \cdot 1000 \cdot 11 = 10890$ for the first row, since there are 11 distinct possible combinations of parameter values excluding combinations where ρ differs from baseline. The actual average value was 10888.

Turning finally to the results, the mean estimate of the coefficient on r_j term in row 1 is 0.01, with a mean standard error of .11, so the interest rate term is not remotely statistically significant in the typical regression. However, most of the Monte Carlo regressions

¹³Groups for which ρ differs from baseline are excluded because the goal in these equations is to estimate ρ ; the question of what the ‘right’ value of ρ is is complicated if ρ differs across groups.

Row	Sample	Instruments	r_j		$(\Delta \log Y_{t+1})_j$		Average NOBS
			Mean	[.05-.95] Range	Mean	[.05-.95] Range	
1	All But ρ	RDUM	0.01 0.11	[-0.11,0.18] [0.10,0.11]			10888
2	BASE + R	RDUM	0.01 0.10	[-0.11,0.18] [0.10,0.10]			2961
3	All But ρ	RDUM +GDUM	-0.02 0.12	[-0.17,0.17] [0.10,0.18]	0.98 0.14	[0.82,1.27] [0.08,0.27]	10888
4	BASE+ R + G	RDUM +GDUM	-0.02 0.12	[-0.17,0.18] [0.10,0.16]	0.98 0.14	[0.82,1.20] [0.08,0.23]	4944
5	All But ρ	RDUM + V_t	0.01 0.09	[-0.10,0.15] [0.09,0.09]	0.11 0.01	[0.09,0.12] [0.01,0.01]	10888
6	BASE+ R	RDUM + V_t	0.01 0.09	[-0.10,0.14] [0.08,0.09]	0.11 0.02	[0.08,0.13] [0.01,0.02]	2961

Notes: The first column indicates which simulated consumers are in the sample. For example 'All but ρ ' means that all simulated consumers are included except those for whom ρ differs from its baseline value.

The second column indicates which categories of dummy variables are used as instruments. For example RDUM indicates use of three dummy variables indicating which of the three possible interest rates the consumer faces.

Table 2: Log-Linearized Euler Equation Estimated on Simulated Data

would be able to reject the true value of $1/\rho = 1/3$ with a high degree of confidence. Results are similar in row 2 when the sample is restricted to the BASE and R groups only. Thus, estimation of the standard log-linearized Euler equation for consumption does not reveal the intertemporal elasticity of substitution *even for consumers behaving exactly according to the model*.

The next row of table 2 presents the results when the basic log-linearized Euler equation is augmented with a term reflecting the predictable growth rate of income, as in equation (11), and income growth is instrumented using the set of dummy variables GDUM, which indicate which permanent-income-growth group the consumer belongs to (RDUM remains in the instrument set to instrument for the interest rate). Again the equation is estimated for two samples, one which includes members with all appropriate parametric combinations, and one containing only consumers who are members of the R and G groups. In row 3, the mean coefficient on the predictable growth rate of income is 0.98, highly significantly different from zero, but not significantly different from one. Results are similar in row 4, which again restricts the sample to the set of consumers for whom one might expect the best results. Furthermore, in the typical regression the coefficient

on the interest rate term is again not significantly different from zero. This result, consumption growth equal to predictable permanent income growth but independent of the interest rate, is precisely what the analysis in Section 3 and in Carroll (1996, 1997) showed holds if consumers are distributed according to the ergodic distribution. Apparently, at least under the parameter values considered here, 9 years of presample simulation for 1000 consumers suffice to generate a sample that generates behavior very similar to that under the ergodic distribution.

As noted in the literature survey above, empirical point estimates of the excess sensitivity of consumption growth to predictable income growth have mostly fallen in the range from 0.0 to about 0.6. Although many of the studies could not reject a coefficient of 1 on the income growth term, it is not possible to claim that the empirical evidence is *more* consistent with a coefficient of 1 than with a coefficient of 0. It might seem, then, that these results rescue the Euler equation from the Scylla of a prediction that $\alpha_1 = 0$ only to smash against the Charybdis of a prediction that $\alpha_1 = 1$. Fortunately, this is not the case. The theoretical arguments and simulation evidence presented do not necessarily imply a coefficient of 1 on $E_t \Delta \log Y_{t+1}$ – they imply a coefficient of one on $E_t \Delta \log P_{t+1}$. That is, consumption should on average grow at the rate of *permanent* income growth. None of the theoretical or simulation work up to this point in the paper has indicated what the coefficient should be on predictable *transitory* movements in income.

The last two rows of the table present the model’s predictions about the coefficient on the predictable transitory movements in income. (Transitory movements in income are predictable because the level of the transitory shock is white noise. Thus, if income is temporarily low today, income growth between today and tomorrow is likely to be high, and vice versa. Hence the instrument used for $E_t \Delta \log Y_{t+1}$ is V_t .) The coefficient on transitory movements in income is statistically significantly different from zero, but, at around .10, is much closer to zero than to one. As before, the coefficient on the interest rate term is insignificantly different from zero.

These very different results for transitory and for permanent income growth imply that there is little we can say about the model’s prediction for the coefficient on predictable

income growth, if we have not also decomposed that growth into the part representing transitory growth and the part representing permanent growth. Essentially all we can say is that (under this set of baseline parameter values), the coefficient on predictable income growth should be somewhere between .10 and 1.0. Of the roughly twenty studies cited by Browning and Lusardi (1996), none (to my knowledge) attempts to decompose predictable income growth into predictable transitory and predictable permanent components. Since standard confidence intervals for α_1 in these papers always overlap the range between 0.10 and 1.0, if ‘excess sensitivity’ is defined as a degree of sensitivity inconsistent with unconstrained intertemporal optimization, *none of the ‘excess sensitivity’ tests summarized by Browning and Lusardi (1996) provides any evidence on whether consumption actually exhibits excess sensitivity to predictable changes in income.*

These results also bear on the finding of Campbell and Mankiw (1989) that regressions aggregate consumption growth on predictable aggregate income growth find a coefficient of roughly 0.5. Although Campbell and Mankiw interpreted their findings as suggesting that about half of consumers behave according to a ‘rule-of-thumb’ and set their consumption equal to their income, Campbell and Mankiw did not decompose their predictable income growth term into a predictable permanent growth term and a predictable transitory term, so it is quite possible that their results are consistent with the standard model without the need for introducing ‘rule-of-thumb’ consumers.

A final category of tests should be mentioned briefly: empirical estimates of the rate of time preference. Lawrance (1991) estimates an equation like (11) using data from the PSID, but including dummy variables for education in the estimating equation. She finds that consumers with more education have higher rates of consumption growth. She concludes that consumers with more education must be more patient. This conclusion would be warranted if the log-linearized consumption Euler equation were valid, because $-\rho^{-1}\delta_j$ is omitted from the baseline empirical specification since δ_j is unobserved. However, given that a positive correlation between permanent income growth and education is a bedrock empirical result in labor economics, an obvious alternative explanation of Lawrance’s results is that the higher consumption growth for more educated consumers reflects their

faster predictable permanent income growth, not a greater degree of patience.

To summarize, when the log-linearized consumption Euler equation is estimated on household data generated by consumers behaving exactly according to the standard model, using the methods that have been used by most of the existing cross-section empirical studies, the results provide no information on either the coefficient of relative risk aversion or on whether consumption exhibits ‘excess sensitivity’ to predictable income growth.

5.2 The Second Order Approximation

A few empirical studies, of which Dynan (1993) is one of the earliest and best, have avoided the log-linearized Euler equation and instead used the second-order approximation to the Euler equation, equation (8),

$$\Delta \log C_{t+1} \approx \rho^{-1}(r - \delta) + \frac{1 + \rho}{2} \eta_{t+1}^2 \quad (15)$$

as the basis of their empirical estimation, using an estimating equation of the form

$$(\Delta \log C_{t+1})_j = \alpha_0 + \alpha_1 r_j + \alpha_2 (\eta_{t+1}^2)_j. \quad (16)$$

where the understanding has been that that the estimation should yield $\alpha_0 = \rho^{-1}\delta$, $\alpha_1 = \rho^{-1}$, and $\alpha_2 = \frac{1+\rho}{2}$. There is a widespread impression that, if any instruments can be found that are correlated with $(\eta_{t+1}^2)_j$, estimation of this equation gets around whatever problems there may be with the log-linearized Euler equation.

Unfortunately, the situation is much subtler than it appears. Obtaining consistent estimates for α_1 and α_2 requires instruments that can identify independent variation in r_j and $(\eta_{t+1}^2)_j$. But recall that according to equation (14)

$$(\eta_{t+1}^2)_j \approx \frac{2}{1 + \rho_j} [g_j - \rho_j^{-1}(r_j - \delta_j)]. \quad (17)$$

Assuming ρ_j is constant across groups and that the second-order approximation is valid, this equation tells us that any instrument correlated with η_{t+1}^2 must be providing information about either r_j, δ_j or g_j . Note, however, that an instrument correlated with r_j is not useful in estimating α_2 , because the variation in η_{t+1}^2 due to variations in r_j will obviously be perfectly correlated with the direct variation in r_j , whose coefficient, remember, is already being estimated by α_1 .

One might hope that an instrument correlated with δ_j could serve to identify α_2 . Certainly, an instrument correlated with δ_j should generate variation in target wealth and therefore in η_{t+1}^2 and so may look useful in the first-stage IV tests. And it is quite plausible to suppose that the time preference rate is correlated with observable variables such as, say, education (one of the instruments Dynan used for η_{t+1}^2). The first row of table 3 therefore presents the results when equation (16) is estimated on simulation data using dummy variables for the time preference rate and interest rate as instruments for η_{t+1}^2 . The coefficients on both the interest rate term and the η_{t+1}^2 term are insignificantly different from zero - just as in Dynan's (1993) empirical work. Note that, if there were *not* econometric problems of some sort, a coefficient of zero on r_j would imply $\rho = \infty$, while a coefficient of zero on η_{t+1}^2 would imply $\rho = -1$, making nonsense of the model.

The econometric problem, of course, is that δ_j also enters the Euler equation in another place: in the $\rho^{-1}(r_j - \delta_j)$ term. But this means that δ_j is an unobserved variable that is correlated with the included variable η_{t+1}^2 , a situation that implies that the coefficient estimate on η_{t+1}^2 will be inconsistent. This example illustrates the point that no instrument that is correlated with the time preference rate will be valid, even if it works well in the first-stage regressions. Furthermore, a test of overidentifying restrictions (such as Dynan performs) will not detect this problem because OI tests only find correlations of instruments with the dependent variable which are *not* captured by the variables that are included, but since $(\eta_{t+1}^2)_j$ is included the OI test should not reject the specification.

A simple thought experiment may clarify the problem better than the foregoing analysis. Consider attempting to estimate equation (8) using data from several groups of consumers who differ from each other in their (observable) interest rates and in their (unobservable) time preference rates, but who have identical g 's. The r_j and $(\eta_{t+1}^2)_j$ terms will vary across groups; first-stage IV regressions of η_{t+1}^2 on the instruments will find the instruments have significant predictive power. Yet the analysis above showed that each of these groups should have consumption growth on average equal to their permanent income growth - that is, all the groups will have identical consumption growth. Hence the coefficient estimates on both r_j and $(\eta_{t+1}^2)_j$ will be zero.

Row	Sample	Instruments	r_j		$(\eta_{t+1}^2)_j$		Average NOBS
			Mean	[.05-.95] Range	Mean	[.05-.95] Range	
1	All	RDUM+	0.00	[-0.17,0.18]	0.18	[-3.19,2.57]	10897
		DELDUM	0.11	[0.09,0.15]	1.94	[1.10,3.47]	
2	BASE + R	RDUM+	0.07	[-0.13,0.25]	9.67	[7.95,12.05]	4957
		GDUM	0.13	[0.11,0.16]	1.45	[0.98,2.21]	
3	All	All But RHODUM	0.00	[-0.14,0.18]	-0.05	[-0.60,0.47]	7931
		and DELDUM	0.11	[0.10,0.11]	0.30	[0.27,0.34]	

Notes: The first column indicates which simulated consumers are in the sample. For example 'All but ρ ' means that all simulated consumers are included except those for whom ρ differs from its baseline value.

The second column indicates which categories of dummy variables are used as instruments. For example RDUM indicates use of three dummy variables indicating which of the three possible interest rates the consumer faces.

Table 3: Second-Order Approximation Estimated on Simulated Data

The conclusion is that, because $(\eta_{t+1}^2)_j$ is a function only of r , g , δ , and ρ and because δ and ρ are unobservable, equation (8) can only be estimated consistently, even in principle, by using a set of instruments that 1) contain independent information on r_j and g_j , and 2) are uncorrelated with preferences. As a practical matter, it is likely to be hard to plausibly identify such instruments, but there is of course no difficulty in simulated data. The next row of table 3 presents the results when equation (8) is estimated on a simulated dataset that should represent the ideal set of circumstances for estimating such an equation: The only differences among the consumers included in this dataset are in r_j and g_j , where r_j is directly observed and g_j is indirectly observed via the set of dummy variables indicating which of three growth-rate groups the consumer belongs to.

The results are interesting. While the coefficient on the interest rate term is still insignificant, the mean coefficient on the η_{t+1}^2 term is 9.7; since equation (8) implies that this coefficient is equal to $\frac{1+\rho}{2}$, this would imply a coefficient of relative risk aversion of over 18. With the mean standard error estimated at about 1.45, the typical regression in this sample would be able to reject the (true) value of $\rho = 3$ with an overwhelming degree of statistical significance.

Why does estimation of this equation fail? Recall the two critical assumptions used in deriving the expression for η_{t+1}^2 upon which the entire foregoing analysis rests. The first was that consumers in each of the j groups were distributed according to an ergodic

distribution which they are assumed eventually to reach. The earlier simulation results showing that average consumption growth is essentially equal to average permanent income growth, and the figure showing that the distribution of x after 10 periods is virtually identical to the steady-state distribution, suggested that this assumption is probably reasonable. The problem therefore must lie in the second assumption: that the *second-order* approximation to the Euler equation is sufficient to capture the important nonlinearities in the problem.

Another way of putting this is to say that the results indicate that the $E_t\eta_{t+1}^2$ term is correlated with higher-order terms in the Taylor expansion of the true function, because if η_{t+1}^2 were *not* correlated with higher-order terms then the coefficient estimate on $E_t\eta_{t+1}^2$ should be unbiased.

The fact that there are missing higher-order terms in equation (8) also undermines the conclusion that $(\eta_{t+1}^2)_j$ is a function only of r , g , δ , and ρ . In particular, there is no longer any reason to exclude the possibility that $(\eta_{t+1}^2)_j$ could be correlated with, for example, the variances of the innovations to transitory and permanent income, $(\sigma_n^2)_j$ and $(\sigma_v^2)_j$. The last regression in table 3 therefore presents the results when the instrument set is expanded to include the dummy indicator variables for σ_n^2 and σ_v^2 . The effect is dramatic: the coefficient on η_{t+1}^2 becomes -0.05, and is no longer significantly different from zero – again reproducing Dynan’s result.

In sum, IV estimation of the second-order approximation to the consumption Euler equation fares little better than IV estimation of the log-linearized equation. Neither approach appears capable of identifying structural parameters even in a dataset consisting exclusively of consumers behaving exactly according to the model.

6 What Is To Be Done?

IV estimation of approximated Euler equations estimation has been a mainstay of economic analysis of consumption for a long time. If the argument of this paper is accepted, such estimation will be abandoned. What kinds of analysis can replace it?

Row	Problems	Mean Estimates	[.05-.95] Range
1	None	3.38 0.61	[2.36,4.65] [0.39,0.89]
2	Signal/Noise = 1/2	2.18 0.36	[1.67,2.94] [0.28,0.45]
3	Signal/Noise = 1/3	1.42 0.23	[1.09,1.96] [0.19,0.30]
4	Assumed $\beta = 1$	2.82 0.79	[1.87,4.34] [0.64,1.05]
5	Assumed $\beta = 1/1.08$	3.79 0.60	[2.85,4.87] [0.42,0.86]

Notes: Results summarize 100 Monte Carlo simulations. First column indicates the problems GMM estimation might face.

Table 4: Euler Equation Estimated on Simulated Data Using GMM

6.1 Bad Ideas

6.1.1 GMM Estimation

The obvious answer is that, since approximation error is the root of all the evils described above, the solution is to dispense with approximation by estimating the full nonlinear Euler equation using the Generalized Method of Moments methodology introduced by Hansen (1982). Table 4 presents the results of GMM estimation on the baseline set of simulated consumers. As expected, the Monte Carlo results imply that GMM estimation usually produces an estimate of the coefficient of relative risk aversion that is not significantly different from the true value $\rho = 3$.

The problem with full-fledged GMM estimation is that consistent estimation requires perfect data on consumption, whereas the available consumption data for households are almost certainly very noisy. Shapiro (1984) estimates that 92 percent of the variation in the PSID food consumption variable is noise; Runkle (1991) estimates that 76 percent of the variation is noise. And although Dynan does not estimate the noise-to-signal ratio in her quarterly *Consumer Expenditure Survey* data, she reports that the standard deviation of quarterly changes in log consumption is 0.2, which seems far too large to reflect quarterly reevaluations of the sustainable level of consumption.

The effect of measurement error on the GMM estimates is illustrated in the second and third rows of table 4. Row 2 reflects the results when the same data on C_{t+1}/C_t that are used for row 1 are first multiplied by a mean-one white noise shock whose distribution is identical to that of the consumption shock. This distributional assumption is motivated by its implication that the signal-to-noise ratio in the resulting data is exactly $1/2$, as indicated in the second column of the table. When GMM is performed on the mismeasured data, the mean estimate of ρ is about 2.2, with an estimated standard deviation of .36, so a hypothesis test that $\rho = 3$ would almost always reject. Row 3 shows that when the signal/noise ratio is reduced to $1/3$ (by multiplying by another white noise shock constructed along the same lines as the first one), the estimate of ρ drops to about 1.4, and the standard error falls.

Another problem with GMM estimation is that estimation of ρ requires an assumption about R and β (or, if R is observed, at least an assumption about β). The last two rows of the table present the results that emerge if the econometrician falsely assumes that $\beta = 1$ (row 4) or $\beta = 1/(1.08)$ (row 5).¹⁴ Assuming that consumers are more patient than the truth reduces the mean estimate of ρ by about 0.6, while assuming that they are less patient boosts the estimated ρ by about 0.4. These results suggest that this problem is less serious than the problems caused by measurement error.

Despite these results, GMM estimation is not completely useless: Because measurement error should bias the estimate of the coefficient of relative risk aversion downward, and because mistaken assumptions about $R\beta$ do not distort the estimates of ρ too badly, the GMM estimate can serve as a rough lower bound on the coefficient of relative risk aversion. A finding of a relatively large lower bound (say, two) would provide moderately interesting information about preferences.

¹⁴All that matters for these equations is the product $R\beta$, so separate experiments showing the results for incorrect assumptions about R would be redundant.

Row	Sample	r_j		g_j		Average NOBS
		Mean	[.05-.95] Range	Mean	[.05-.95] Range	
1	BASE + RDUM + GDUM	-0.01 0.02	[-0.04,0.02] [0.01,0.02]	0.10 0.02	[0.07,0.13] [0.01,0.03]	4954

Notes: The first column indicates that the sample consists only of the consumers with baseline parameter values or those for whom either the interest rate or the growth rate of income differs from baseline.

Table 5: Regression of $(\eta_{t+1}^2)_j$ On r_j and g_j

6.1.2 Using η_{t+1}^2 As the Dependent Variable

Equation (14), reproduced below for convenience, appears to offer hope of estimating the coefficient of relative risk aversion even without GMM estimation:

$$(\eta_{t+1}^2)_j \approx \frac{2}{1+\rho} [g_j - \rho_j^{-1}(r_j - \delta_j)]. \quad (18)$$

In principle, one could estimate this equation using data from groups of consumers with different values of g and r , so long as there were no differences in δ or ρ across those groups. If the second order approximation were good, the coefficient on g should equal $2/(1+\rho)$ and that on r should equal $(2/\rho(1+\rho))$.

Table 5 presents the results when this equation is estimated using the best possible subset of consumers from the simulated dataset.¹⁵ The estimated coefficient on g is about 0.10 and is highly statistically significant. But this estimate implies an estimate of $\rho = 2/.10 - 1 \approx 19$. Since the true ρ is 3, this is not an attractive result. The coefficient on the interest rate is approximately zero and statistically insignificant. Thus, equation (18) also fails to provide a consistent way to estimate ρ .

6.1.3 Individual-Specific Euler Equation Estimation

The arguments to this point in the paper have been directed at demonstrating that the traditional Instrumental Variables approach to Euler equation estimation does not succeed. Because all RHS variables were always instrumented with group identifiers, the second-stage regressions contained no individual-specific information in the independent variables.

¹⁵The only consumers included were those from the baseline group and those groups for whom r or g varied from the baseline.

For example, each individual's idiosyncratic expectation of $\eta_{i,t+1}^2$ was effectively replaced by the mean value of $\eta_{i,t+1}^2$ for the group to which that consumer belonged.

The logic proposed as an explanation for the failure of the estimation method relied on the proposition (verified by simulations) that the group mean values of the $\eta_{i,t+1}^2$ terms would take on particular values. That logic, therefore, does not necessarily prove that it is impossible to estimate structural consumption Euler equations using idiosyncratic, individual-specific data. If it were possible to observe, for each individual i , their idiosyncratic, contemporaneous value of $E_{i,t}\eta_{i,t+1}^2$, then it might be possible to estimate equation (8) *without* using instrumental variables. To be specific, one could estimate:

$$\Delta \log C_{i,t+1} = \alpha_0 + \alpha_1 r_i + \alpha_2 E_{i,t} \eta_{i,t+1}^2 + \epsilon_{i,t+1}. \quad (19)$$

Paxson and Ludvigson (1997) and Laibson (1997) both examine this possibility theoretically. Paxson and Ludvigson are able to calculate, at each possible wealth state, the approximation bias in the estimate of ρ associated with the second-order approximation. Although they are only able to solve their model back to six periods before the end of life, they find that the bias is substantial at most levels of wealth; the biased estimate of ρ is typically around 0.7 or 0.8 times the true ρ , although the extent of the bias declines as wealth rises.

Laibson (1997) adopts a methodology somewhat more akin to that of this paper. He solves a lifetime optimization problem, then simulates a population of consumers behaving exactly according to the model. For each simulated consumer at each age, he calculates the model's mathematical expectation for $E_{i,t}\eta_{i,t+1}^2$, then performs a regression like that of equation (19) on the simulated data (omitting the r_i term because he does not allow variation in interest rates across households.) Like Paxson and Ludvigson (1997), he finds that the resulting estimate of ρ is downward biased by a factor of around 0.8.

Table 6 presents the results when the corresponding experiment is performed in my model under the baseline set of parameter values, and under several alternative parametric configurations. Under the baseline parameter values, the point estimate of α_2 is 4.46, which implies an estimate of about $\rho = 8$ — an *upward* bias, in contrast to the Laibson/Paxson-

Row	Sample	Method	$E_t \eta_{t+1}^2$		Average NOBS
			Mean	[.05-.95] Range	
1	BASE	OLS	4.46 0.75	[3.03 6.38] [0.55 0.93]	991
2	R=1.00	OLS	4.50 0.67	[3.34 5.92] [0.49 0.81]	996
3	R=1.04	OLS	4.26 0.78	[2.76 6.06] [0.56 0.99]	987
4	$\beta = 1.00$	OLS	4.16 1.01	[2.52 6.50] [0.73 1.26]	991
5	$\beta = 1/1.08$	OLS	4.78 0.65	[3.63 6.01] [0.56 0.74]	997
6	$g = .02$	OLS	5.88 1.50	[3.27 9.44] [1.03 1.87]	993
7	$g = .06$	OLS	3.94 0.50	[2.72 5.13] [0.38 0.58]	990
8	$\rho = 1$	OLS	4.27 0.48	[3.55 4.99] [0.41 0.54]	989
9	$\rho = 5$	OLS	5.52 0.93	[3.69 7.41] [0.74 1.15]	993
10	$\sigma_v = .05$	OLS	5.01 1.15	[3.28 7.06] [0.88 1.52]	986
11	$\sigma_v = .15$	OLS	3.71 0.47	[2.79 4.87] [0.37 0.59]	995
12	$\sigma_n = .05$	OLS	5.16 0.44	[3.54 6.55] [0.31 0.56]	998
13	$\sigma_n = .15$	OLS	4.24 1.50	[1.88 6.72] [1.18 1.78]	994
14	BASE	IV	8.05 1.58	[5.62 11.25] [0.80 2.90]	991
15	R=1.00	IV	7.75 1.28	[5.95 10.87] [0.74 2.41]	996
16	R=1.04	IV	7.88 1.72	[4.98 12.17] [0.76 3.18]	987

Notes: Sample column indicates which parts of the simulated data are used.
Method column indicates whether OLS or IV is used.

Table 6: Second-Order Approximation Using Idiosyncratic Data

Ludvigson findings of downward biases.¹⁶

If it were possible to be confident about the exact magnitude of the bias in the estimate of ρ using this method, it might be at least remotely possible to obtain a reliable estimate of the value of ρ by estimating an equation like (19) and then correcting for bias. However, rows 2-13 of table 6 show that when the same estimation exercise is performed on each of the other groups, the magnitude of the bias is somewhat affected by the value of the other parameters in the model, both observable and unobservable. Without reliable independent information on these parameters (particularly the taste parameters) at the individual level, it is not possible to know the exact magnitude of the bias.

As a way of investigating the source of this bias, Figure 3 plots the true numerical expectation of $E_t \Delta \log C_{t+1}$ as a function of the level of cash-on-hand under the baseline parameter values, along with the expected value of the second-order approximation (8). The minimum and maximum values of x_t for the plot are the first and 99th percentiles in the ergodic distribution of x_t that arises from the simulations. The figure shows that the second order approximation does a remarkably poor job capturing the relationship between cash-on-hand and expected consumption growth over the range of values of x_t that arise during the simulations. However, it is easy to see from this figure why the coefficient estimates on η_{t+1}^2 are biased upward: as wealth gets lower and lower (and therefore η_{t+1}^2 gets larger and larger), the second-order approximation falls further and further below the true value of expected consumption growth. Since, in the regressions, the coefficient on $E_t \eta_{t+1}^2$ is not constrained to be $\frac{\rho+1}{2}$, the regression chooses a much larger value for that coefficient, with an offsetting adjustment to the intercept to get the mean level of the function right.

Of course, in principle a high-enough order approximation to the Euler equation could capture the expected consumption growth function arbitrarily well. However, figure 3 shows that even a fourth-order approximation does not do a very good job of capturing the relationship between consumption growth and cash-on-hand. Given the limitations of

¹⁶I do not know why Laibson and Paxson and Ludvigson obtained downward biases. One possibility is that their data may have been dominated by consumers who were effectively patient, while my consumers are all impatient.

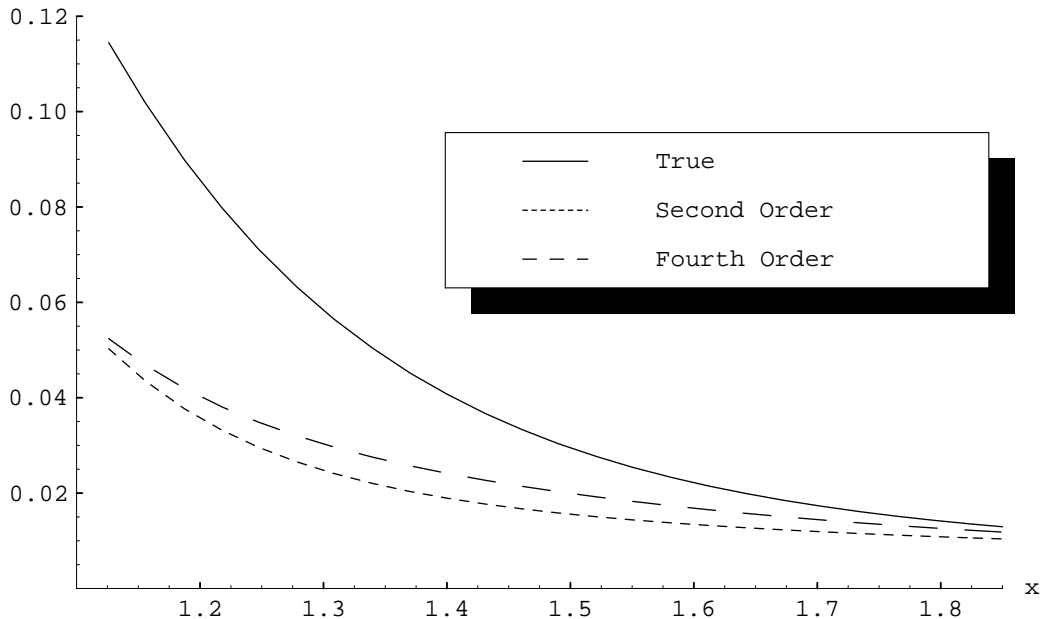


Figure 3: True and Approximated $E_t \Delta \log C_{t+1}$

actual data, it seems clear that it will not be possible to estimate the coefficient of relative risk aversion with much precision using any plausible approximation to the consumption Euler equation.

6.2 Good Ideas

6.2.1 Consumption Growth Regressions

It is important to make a distinction between estimating Euler equations and estimating regressions of consumption growth on explanatory variables. Euler's name is implicated in the standard terminology as shorthand for the idea that one is estimating a first-order condition from a maximization problem. While I believe that the arguments of this paper demonstrate the impossibility of recovering a direct estimate of structural parameters from consumption growth regressions, there are nevertheless several kinds of consumption growth regressions that could be used to test important implications of models of intertemporal optimization. Two such tests have already been implicitly suggested. Table 2 showed that, under configurations of parameter values that generate buffer-stock saving, a regression of consumption growth on the predictable component of permanent

income growth should yield a coefficient near one, while the coefficient on the predictable component of transitory income growth should be much smaller (around .10 for baseline parameter values). These are eminently testable propositions.¹⁷

Given the results of Table 6, it even seems worthwhile to attempt to estimate an equation of the *form* of the second-order approximation to the Euler equation (but only if idiosyncratic data are used). The point of the earlier discussion of Table 6 was that the coefficient on $E_{i,t}\eta_{i,t+1}^2$ did not yield an unbiased estimate of ρ . From a less structural point of view, however, the lesson of the table is that for any tested set of parameter values the model implies a hugely statistically significant relationship between consumption growth and $E_{i,t}\eta_{i,t+1}^2$.

Of course, as a practical matter, an econometrician never observes each household's idiosyncratic expectations of a variable like $\eta_{i,t+1}^2$, so the research strategy just described cannot be implemented directly. However, in the theoretical model, $E_{i,t}\eta_{i,t+1}^2$ is a monotonic function of cash-on-hand $x_{i,t}$, which *is* observable. This suggests that it should be possible to estimate the equation using $x_{i,t}$ (and perhaps higher moments of x) as instruments for $\eta_{i,t+1}^2$. Row 14 of table 6 presents the results when the equation is estimated using $x_{i,t}$ and $x_{i,t}^2$ as instruments for $\eta_{i,t+1}^2$. The coefficient estimate on the instrumented $\eta_{i,t+1}^2$ term remains highly statistically significant, and is even larger than the value that it takes when the equation is estimated using the individual-specific values of $E_{i,t}\eta_{i,t+1}^2$ taken from the model. Rows 15 and 16 show that similar results obtain for two of the other groups of consumers; for brevity, results for the remaining groups are omitted. These last three regressions are feasible in many if not most of the datasets that have been used to estimate the traditional consumption Euler equation in the past. Estimating such an equation would be a particularly easy task for any author who has estimated a traditional Euler equation in one of these datasets and still has the computer code available.

¹⁷A hint of the answer, at least for predictable permanent growth, is already available: the work by Carroll and Summers (1991) showing that consumption growth parallels income growth over most of the working lifetime strongly suggests that when the experiment is performed properly the coefficient on predictable low-frequency growth in income will be close to one; the results in Carroll (1994) also support such an interpretation.

6.2.2 Other Ideas

Another particularly promising avenue is to test the model's predictions about the determinants of target or buffer-stock wealth. Table 7 presents the results when the level of wealth is regressed on the set of variables that are, in principle, observable at either the individual level or the group level. The effects are all in the directions one would expect: higher interest rates encourage more wealth-holding; higher permanent income growth depresses wealth through standard human wealth channels; consumers facing higher interest rates hold more wealth; consumers facing greater income uncertainty also hold more wealth; and consumers who are more risk averse hold more wealth.¹⁸ Note that several of these variables have very high degrees of statistical significance in the typical regression. To my knowledge, the only empirical tests thus far performed along these lines are in Carroll and Samwick (1997), who find, using the PSID, that the variance of both the transitory and permanent shocks to income are positively and significantly related to wealth; and Carroll and Weil (1997), who find a *positive* association between income growth and saving, which they note is inconsistent with a buffer-stock model of saving.

In principle, it is even possible to estimate structural parameter values. A simple example of how this can be done can be found in Carroll and Samwick (1997). Using data from the PSID, they estimate a regression of household wealth on the variance of permanent income shocks. Then, using a buffer-stock model similar to that used in this paper, they determine the value of the rate of time preference such that, if similar regressions were estimated in simulated data from the model, the coefficient estimates would be similar to those obtained from the empirical work. This is a very simple example of a literature on estimation by simulation; for a much more sophisticated example in a different context, see Michaelides and Ng (1997).

Carroll and Samwick (1997) fixed all parameter values but one, and obtained only a

¹⁸It might seem surprising to list the coefficient of relative risk aversion among the observable variables. However, two large survey datasets (the HRS and the PSID) have recently added questions explicitly designed to elicit information about risk aversion. Kimball et al. (1997) report that these variables have some plausible correlations with other observable variables. For example, consumers who report a high degree of risk aversion are less likely to smoke. It would be very interesting to see if such households also hold more wealth, *ceteris paribus*.

Row	Independent Variable	Coefficient Estimate		NOBS	\overline{R}^2
		Mean	[.05-.95] Range		
1	R	1.80 0.17	[1.58,2.05] [0.16,0.17]	2969	0.04
2	G	-7.29 0.19	[-7.62,-6.98] [0.19,0.20]	2970	0.33
3	σ_n^2	14.48 0.44	[13.79,15.28] [0.42,0.45]	2978	0.27
4	σ_v^2	3.00 0.36	[2.38,3.63] [0.34,0.37]	2966	0.02
5	ρ	0.23 0.00	[0.222,0.228] [0.002,0.002]	2969	0.82

Notes: Assumption is that actual values of all variables are directly observed. IV estimation is also possible and should produce consistent estimates.

Table 7: Regressions of Cash-On-Hand On Observable Variables

point estimate for that parameter. An even more ambitious project is to estimate several parameters at once, in such a way that standard errors can also be obtained. Although the technical and computational challenges are formidable, two recent papers have scored impressive success in doing this. Parker and Gourinchas (1996) develop routines to quickly solve and simulate a dynamic life cycle simulation model under arbitrary values of the coefficient of relative risk aversion and the time preference rate. They then use an econometric hill-climbing routine to search for the (ρ, δ) combination that causes their model to best match data from the U.S. Consumer Expenditure Surveys. They obtain plausible and tight parameter estimates for both ρ and δ . And Palumbo (1997) estimates a structural model for precautionary saving for out-of-pocket medical expenditures by the elderly.

In sum, there are many possible avenues for testing models of intertemporal consumption choice even if structural Euler equation estimation must be abandoned.

7 Conclusions

This paper argues that the estimation of consumption Euler equations using instrumental variables methods on cross-section household data should be abandoned because it does not yield any useful information. However, there are many other promising ways to

test models of consumption under uncertainty, and even some ways to get estimates of structural parameters; presumably inventive researchers can come up with many more ways of testing the model.

References

- ATTANASIO, O. P., AND G. WEBER (1985): "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103, 1121–57.
- BARSKY, R. B., F. T. JUSTER, M. S. KIMBALL, AND M. D. SHAPIRO (1997): "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Survey," *Quarterly Journal of Economics*, CXII(2), 537–580.
- BROWNING, MARTIN J., AND ANNAMARIA LUSARDI (1996): "Household Saving: Micro Theories and Micro Facts," *Forthcoming, Journal of Economic Literature*.
- CAMPBELL, J. Y., AND N. G. MANKIW (1989): "Consumption, Income, and Interest Rates: Reinterpreting the Time-Series Evidence," in *NBER Macroeconomics Annual, 1989*, ed. by O. J. Blanchard, and S. Fischer, pp. 185–216. MIT Press, Cambridge, MA.
- CARROLL, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 1992(2), 61–156.
- (1994): "How Does Future Income Affect Current Consumption?," *The Quarterly Journal of Economics*, CIX(1).
- (1996): "Buffer Stock Saving: Some Theory," *Manuscript, Johns Hopkins University, January*.
- (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, CXII(1), 1–56.
- CARROLL, C. D., AND M. S. KIMBALL (1996): "On the Concavity of the Consumption Function," *Econometrica*, 64(4), 981–992.
- CARROLL, C. D., AND A. A. SAMWICK (1997): "The Nature of Precautionary Wealth," *Forthcoming, Journal of Monetary Economics*, 40(1).

- CARROLL, C. D., AND L. H. SUMMERS (1991): "Consumption Growth Parallels Income Growth: Some New Evidence," in *National Saving and Economic Performance*, ed. by B. D. Bernheim, and J. B. Shoven. Chicago University Press, Chicago.
- CARROLL, C. D., AND D. N. WEIL (1997): "Habits and Saving: Theory and Evidence," *Work in Progress*.
- DYNAN, K. E. (1993): "How Prudent Are Consumers?," *Journal of Political Economy*, 101(6), 1104–1113.
- GOURINCHAS, P.-O., AND J. PARKER (1996): "Consumption Over the Life Cycle," *Manuscript, University of Michigan*.
- HALL, R. E. (1978): "Stochastic Implications of the Life-Cycle/Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 96, 971–87.
- HANSEN, L. M. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, pp. 1029–1054.
- LAIBSON, D. (1997): "Hyperbolic Discount Functions and Time Preference Heterogeneity," *Manuscript, Harvard University*.
- LAWRANCE, E. C. (1991): "Poverty and the Rate of Time Preference: Evidence from Panel Data," *Journal of Political Economy*, 99(1), 54–77.
- MICHAELIDES, A., AND S. NG (1997): "Estimating the Rational Expectations Model of Speculative Storage: A Monte Carlo Comparison of Three Simulation Estimators," *Manuscript, Department of Economics, Boston University*.
- PALUMBO, M. G. (1997): "Precautionary Saving and Out-Of-Pocket Medical Expenditures Near the End of the Life Cycle," *Manuscript, University of Houston*.
- PAXSON, C., AND S. LUDVIGSON (1997): "Approximation Bias in Euler Equation Estimation," *Manuscript, Federal Reserve Bank of New York*.

RUNKLE, D. E. (1991): "Liquidity Constraints and the Permanent Income Hypothesis: Evidence from Panel Data," *Journal of Monetary Economics*, 27(1), 39–98.

SHAPIRO, M. D. (1984): "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data," *Economics Letters*, 14(1), 93–100.

ZELDES, S. P. (1989): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 104, 275–298.