

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

Christopher D. Carroll[†]
ccarroll@jhu.edu

Published as: Carroll, Christopher D. (2001) "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)", Advances in Macroeconomics: Vol. 1: No. 1, Article 6.

<http://www.bepress.com/bejm/advances/vol1/iss1/art6>

Keywords: Euler equation, uncertainty, consumption, excess sensitivity

JEL Classification Codes: C6, D91, E21

Abstract

This paper shows that standard methods for estimating log-linearized consumption Euler equations using micro data 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 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 strategies that are not subject to the problems of Euler equation estimation.

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

All of the programs used to generate the results in this paper are available at the author’s website, <http://www.econ.jhu.edu/People/ccarroll/carroll.html>.

[†]Department of Economics, Johns Hopkins University, Baltimore, MD 21218-2685, 410-516-7602 (office), 410-516-7600 (fax).

1 Introduction

Estimation of Euler equations has occupied a central place in consumption research over the more than twenty years since Hall (1978) first derived and tested the consumption Euler equation. Unfortunately, despite scores of careful empirical studies using household data, Euler equation estimation has not fulfilled its early promise to reliably uncover preference parameters like 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 little or no excess sensitivity.

This paper offers an explanation for the conflicting empirical results, by showing that when the Euler equation estimation methods that have been widely used on household data are applied to a set of data generated by simulated consumers behaving exactly according to the standard consumption model, those methods are incapable of producing an econometrically 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.

In principle, the theoretical problems with Euler equation estimation stem from approximation error. The standard procedure has been to estimate a log-linearized, or first-order approximated, version of the Euler equation. This 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 methods for studying 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 solved and simulated. The next section describes the standard empirical methodology for estimating Euler equations and summarizes the results that have been reported in the literature. Section 4 describes the details of the simulations which generate the data to be analyzed. Section 5 is the heart of the paper: It shows that the standard empirical methods cannot produce consistent estimates of true model parameter values. The penultimate section describes several empirical strategies that are candidates to replace Euler equation estimation, and the final section concludes.

2 The Model

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

$$\begin{aligned}
\max_{\{C_t\}} \quad & u(C_t) + E_t \left[\sum_{s=t+1}^T \left(\beta^{s-t} u(\tilde{C}_s) \right) \right] \left(\right. \\
\text{s.t.} \quad & \\
X_{s+1} \quad & = R(X_s - C_s) + Y_{s+1} \\
Y_{s+1} \quad & = P_{s+1} V_{s+1} \\
P_{s+1} \quad & = GP_s N_{s+1} \\
u(C) \quad & = \frac{C^{1-\rho}}{1-\rho} \text{ where } \rho > 1
\end{aligned} \tag{1}$$

where P_s is permanent labor income, which is buffered by lognormally distributed mean-one shocks N with variance of $\log N = \sigma_n^2$, implying that $\log P$ follows a random walk with drift; Y is current labor income, which is equal to permanent labor income multiplied by a mean-one transitory shock V which is equal to zero with probability p (think of this as unemployment) and otherwise is distributed lognormally with variance of $\log V = \sigma_v^2$, and with a mean that guarantees that $E_t[\tilde{V}_{t+1}] = 1$; 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 $\rho > 1$.

The solution to this model obeys the Euler equation:

$$R\beta E_t \left[\tilde{C}_{t+1}/C_t \right]^{-\rho} = 1. \tag{2}$$

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

$$R\beta E_t \left[\frac{\tilde{P}_{t+1} c_{t+1}^*(\tilde{x}_{t+1})}{P_t c_t} \right]^{-\rho} \left(= 1. \tag{3}$$

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

¹The \sim here and henceforth will be used to indicate a variable whose value is uncertain as of the date at which an expectation is being taken.

number of periods from the end of life.² 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 as the horizon recedes, and consumers behaving according to the converged rule can be described as engaging in ‘bu er-stock’ saving. I will denote the optimal consumption rule for any period t as $c_t^*(x_t)$ and the converged consumption rule as $c^*(x) = \lim_{n \rightarrow \infty} c_{T-n}^*(x)$, which corresponds to the infinite horizon solution.

All numerical and simulation results in the paper will be generated from the converged consumption rule. Carroll (1997) argues that empirical evidence for US households suggests that even consumers with finite horizons behave like impatient (but infinite-horizon) ‘bu er-stock’ for much of their working lifetimes; that paper argues that the transition from bu er-stock behavior to something more closely resembling classical life cycle behavior (where the problems emphasized in this paper would be lessened) happens for the median household somewhere between ages 45 and 50. Cagetti’s (1999) recent paper implies a similar age for the transition; Gourinchas and Parker (1999) argue that the transition occurs somewhere around age 40.³ Since most empirical microeconomic work has restricted the sample to households between the ages of 25 and 60 (to avoid including students and others who have not formed a permanent attachment to the labor force on the young end, and early retirees on the older end), under any of these estimates bu er-stock saving behavior should be expected to obtain for a large proportion of the households in the data that has been used in empirical studies.

To verify accuracy of the numerical solution, Figure 1 plots $R\beta E_t[(C_{t+1}/C_t)^{-\rho}]$ as a function of x_t . Errors in the numerical solution will lead the function to differ from one at points away from the gridpoints chosen for x_t (where equality is imposed by the solution method). The figure shows that the errors involved in numerical solution 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. This figure serves to illustrate the point that the problems with Euler equation estimation documented in the rest of the paper are in a sense attributable to the use of approximations to the Euler equation, since (as the figure shows) the true nonlinear Euler equation always holds by construction.

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

²For more details on the method of solution, see Carroll (1992, 1997).

³The difference is probably attributable to the fact that Gourinchas and Parker match mean rather than median behavior. The mean includes many high-income households who save much more than the median household (in proportion to their incomes), and thus could be expected to reflect a higher average degree of patience.

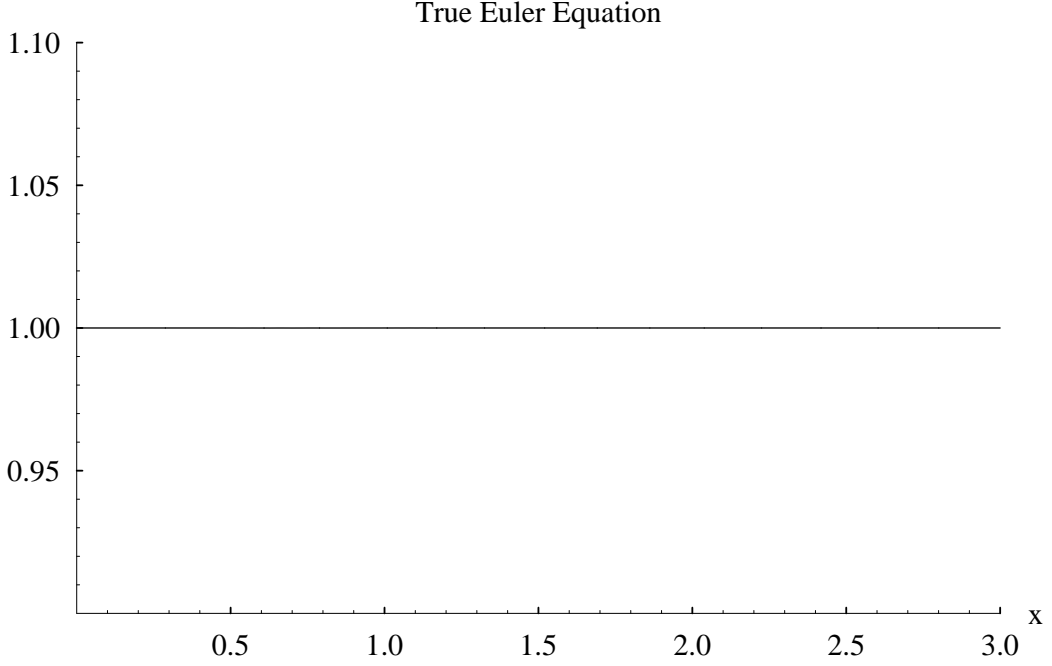


Figure 1: Numerical Value of $R\beta E_t \left[\left(\frac{\tilde{P}_{t+1} c^*(\tilde{x}_{t+1})}{P_t c^*(x_t)} \right)^{-\rho} \right]$

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:

$$R\beta E_t[1 - \rho\tilde{\eta}_{t+1}] \approx 1. \quad (4)$$

A simple transformation of this first-order approximation 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

$$R\beta(1 - \rho E_t[\Delta \log \tilde{C}_{t+1}]) \approx 1. \quad (5)$$

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 $\log(1 - \rho E_t[\Delta \log \tilde{C}_{t+1}]) \approx -\rho E_t[\Delta \log \tilde{C}_{t+1}]$ gives

$$\begin{aligned} (r - \delta) - \rho E_t[\Delta \log \tilde{C}_{t+1}] &\approx 0 \\ E_t[\Delta \log \tilde{C}_{t+1}] &\approx \rho^{-1}(r - \delta), \end{aligned} \quad (6)$$

or, defining the expectation error $\epsilon_{t+1} = \Delta \log C_{t+1} - E_t[\Delta \log \tilde{C}_{t+1}]$, 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 of $(1 + \eta_{t+1})^{-\rho}$ 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

$$\Delta \log C_{t+1} \approx \rho^{-1}(r - \delta) + \left(\frac{\rho + 1}{2}\right) \left(E_t[\tilde{\eta}_{t+1}^2] + \epsilon_{t+1}\right), \quad (8)$$

and if η_{t+1}^2 is uncorrelated with r and δ , then the $E_t[\tilde{\eta}_{t+1}^2]$ term will be absorbed in the constant term of a regression estimate of (7).⁴

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.⁵ The principal purpose of this paper is to show that the usual cross-section procedures for microeconomic estimation of this equation do not work; the penultimate section briefly discusses whether time series estimation methods are similarly problematic.

The instrumental variables approach to estimating the model using microeconomic data can be usefully thought of as equivalent to taking means within groups of consumers with similar characteristics, and identifying parameter values by differences in these group-means. For example, typical instruments used in the empirical literature are education group or occupation group. Henceforth I will denote distinct groups by the subscript j and the group-mean value of a variable X whose value differs across members of the group will be designated $(X)_j$. For example, if we were to designate the growth rate of consumption for an individual household as $\Delta \log C_{i,t+1}$ then the group-mean value of consumption growth across all consumers in group j would be designated $(\Delta \log C_{t+1})_j$ which would be calculated (assuming there are m consumers in group j who happen to have index numbers $i = 1 \dots m$) as $(\Delta \log C_{t+1})_j = (1/m) \sum_{i=1}^m \Delta \log C_{i,t+1}$. Parameters which are assumed to take a common

⁴A 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 $(\rho + 1)/2$.

⁵A 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 (1995).

value for all members of the group are unobtrusively indicated by a subscript j , e.g. ρ_j, r_j , and δ_j . In this notation, 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 has been estimated using regression equations of the form

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

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} . According to equation (9), 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 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 .⁶ However, the poor results in estimating ρ have often been interpreted as reflecting poor identifying information about exogenous differences in r across groups, rather than as important rejections of the Euler equation itself.⁷

The potential empirical problems with identifying exogenous variation in interest rates across households have led many authors to focus on another feature of the model: Hall's 'random walk' proposition. Hall (1978) 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 the change in consumption between t and $t + 1$. The alternative hypothesis has usually been that consumption is 'excessively sensitive' to forecastable income growth. Formally, denoting the expected growth rate of income as $E_t[\Delta \log \tilde{Y}_{t+1}]$, the equation most commonly estimated has been:

$$(\Delta \log C_{t+1})_j = \alpha_0 + \alpha_1 r_j + \alpha_2 (E_t[\Delta \log \tilde{Y}_{t+1}])_j + \epsilon_j, \quad (11)$$

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

Empirical results estimating equation (11) using micro data have been hardly better than those estimating the baseline equation (10).⁸ 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 apologist for the model might note that most estimates are in the range between 0 and 0.6.

⁶See the survey paper by Browning and Lusardi (1996) for more details.

⁷Usually identification has been obtained by calculating a marginal tax rate for each household 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.

⁸Although, interestingly, when the equation is estimated using aggregate 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, fatally, r_j and $(E_t[\Delta \log \tilde{Y}_{t+1}])_j$. Those papers show that ‘impatient’ consumers behaving according to the standard CRRA intertemporal optimization model will engage in ‘bu er-stock’ or target saving behavior,⁹ and that, among a collection of bu er-stock consumers with the same parameter values, if the distribution of x across consumers has converged to its 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, 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 bu er-stock model with a target wealth \bar{w} , then it is impossible for consumption growth to be permanently di erent from underlying income growth. If consumption growth were forever greater than permanent labor income growth, consumption would eventually exceed labor income by an arbitrarily large amount, driving wealth to negative infinity. If consumption growth were permanently less than labor income growth, labor 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 di er permanently from average income growth.¹⁰

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

$$(\Delta \log C_{t+1})_j \approx \rho^{-1}(r_j - \delta_j) + \left(\frac{1 + \rho}{2}\right) \left(E_t[\tilde{\eta}_{t+1}^2]\right)_j. \quad (13)$$

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 across consumers in the group is equal to the average of its expected value. Substituting $(\eta_{t+1}^2)_j$ for $(E_t[\tilde{\eta}_{t+1}^2])_j$ 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

⁹The term ‘impatient’ here and henceforth refers to the condition $R E_t[(G\tilde{N}_{t+1})^{-\rho}] < 1$. Note that, so long as income is growing over time $G > 1$, consumers can be impatient in the required sense even if $\rho = 1$ so that they do not discount future utility at all.

¹⁰For much more careful discussion and arguments, see Carroll (1996, 1997).

must take:

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

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,¹¹ there will be a direct relationship between the level of wealth and the value of $E_t[\tilde{\eta}_{t+1}^2]$. 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 lower 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 $E_t[\tilde{\eta}_{t+1}^2]$. Across steady-states, the higher value of the $E_t[\eta_{t+1}^2]$ term should exactly offset the lower 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).¹²

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 econometric 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 average growth rate of consumption will be equal to g for every group 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 $E_t[\tilde{\eta}_{t+1}^2]$ 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 simulations 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. The next section performs the necessary simulations.

¹¹This is an implication of the concavity of the consumption function proven by Carroll and Kimball (1996).

¹²This 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	0.02	0.04	0.06
ρ	1	3	5
σ_n	0.05	0.10	0.15
σ_v	0.05	0.10	0.15

Table 1: Parameter Values

4 The Simulations

The procedure for generating simulated data from the model is as follows. First, I solve the model for the baseline set of parameter values indicated in Table 1, 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.¹³ 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 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 numerical 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 $\Delta \log C_{t+1}$, r , $\Delta \log Y_{t+1}$, and the dummy variables indicating group membership for each group. With the exception of the interest rate, the simulated 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.

¹³Thus, there are 13 groups altogether: the baseline group plus one positive and one negative deviation from the baseline parameter value for each of six parameters.

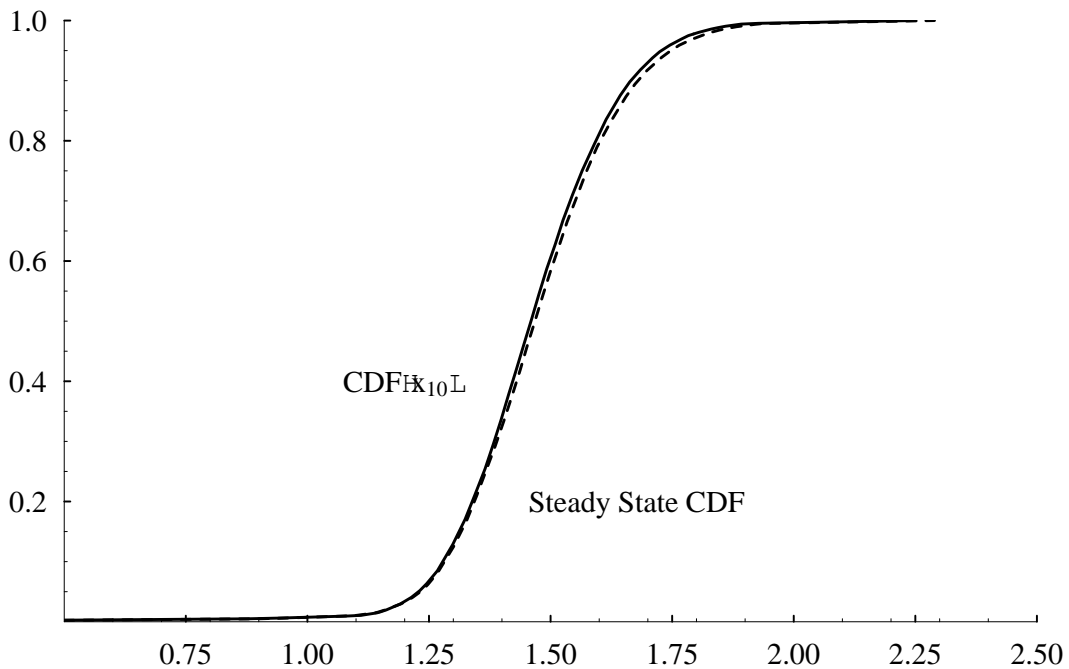


Figure 2: Distribution of x After 10 Years (Solid) vs. Ergodic Distribution (Dashing)

The goal is to characterize the kinds of regression results that an econometrician would obtain using a sample of data drawn from these simulations. 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 10,000 times to obtain a distribution of parameter estimates and standard error estimates.

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 such ideal circumstances, Euler equation estimation by standard microeconomic IV methods does not work. 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 of the six panels presents baseline results when equal numbers of consumers from each possible parametric combination (except for deviations of ρ from baseline) are included.¹⁴ (This sample selection is indicated by the text ‘All but ρ ’ under the ‘Consumers in Sample’ column). 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 regression reported in panel, 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 10,000 times with 10,000 different randomly-chosen collections of 1000 simulated consumers. For each variable, the table presents the mean (across the 10,000 regressions) of the coefficient estimates and the mean of the estimated 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 0.01, 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 of NOBS is almost exactly right, at 10889.

¹⁴Groups for which ρ differs from baseline are excluded because the goal in these equations is to see if the estimation can uncover the ‘right’ estimate of ρ ; the question of what the ‘right’ value of ρ may be is muddled if ρ differs across groups.

Panel	Consumers in Sample [†]	Instruments [‡]	Coe . on r_j		Coe . on $(\Delta \log Y_{t+1})_j$		Average NOBS
			Mean*	[.05-.95] Range	Mean*	[.05-.95] Range	
1	All But ρ	RDUM	0.00 (0.14)	[-0.23,0.24] (0.14,0.14)			10889
2	BASE + R	RDUM	0.00 (0.15)	[-0.23,0.24] (0.15,0.15)			2970
3	All But ρ	RDUM +GDUM	0.00 (0.14)	[-0.20,0.20] (0.11,0.19)	0.97 (0.14)	[0.79,1.20] (0.08,0.25)	10889
4	BASE+ R + G	RDUM +GDUM	0.00 (0.13)	[-0.20,0.20] (0.10,0.18)	0.97 (0.14)	[0.79,1.20] (0.07,0.25)	4950
5	All But ρ	RDUM + V_t	0.00 (0.12)	[-0.21,0.21] (0.12,0.13)	0.09 (0.01)	[0.07,0.10] (0.01,0.01)	10889
6	BASE+ R	RDUM + V_t	0.00 (0.13)	[-0.20,0.21] (0.12,0.14)	0.11 (0.02)	[0.07,0.14] (0.02,0.02)	2970

Notes: The dependent variable in all regressions is the growth rate of consumption, $\Delta \log C_{i,t+1}$.

[†]The column labelled 'Consumers In Sample' indicates which simulated consumers are included the sample. For example 'All but ρ ' in panel 1 means that all simulated consumers are included except those for whom ρ differs from its baseline value, while 'Base + R' in Panel 2 means that the sample includes consumers with baseline parameter values and those for whom R differs from the baseline.

[‡]The column labelled 'Instruments' indicates which categories of dummy variables are used as instruments for r and $(\Delta \log Y_{t+1})$ when the regression is estimated on the data generated by the simulations. For example RDUM indicates use of three dummy variables indicating which of the three possible interest rates the consumer faces.

*The first row in each panel presents the average value and range of the coefficient estimates across the Monte Carlo simulations. The second row in each panel presents the average value and range of values for the regressions' estimates of the standard error.

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

Turning finally to the results, the mean estimate of the coefficient on the r_j term in panel 1 is 0.00, with a mean standard error of 0.14, so the interest rate term is not remotely statistically significant in the typical regression. Furthermore, most of the Monte Carlo regressions would be able to reject the true value of $1/\rho = 1/3$ with a high degree of confidence.

Panel 2 narrows the sample to the set which offers the best hope, in econometric terms: It excludes all consumers who differ from the baseline parameter values in any way *other* than in the interest rate they face. (The only consumers in the sample are the 'BASE' and 'R' groups). Results are virtually identical to those in panel 1. 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 second panel 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 next panel 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 panel 3, the mean coefficient on the predictable growth rate of income is 0.97, highly significantly different from zero, but not significantly different from one. Results are similar in panel 4, which restricts the sample to the set of consumers for whom one might expect the best results for the Euler equation method. Furthermore, in the typical regression the coefficient on the interest rate term is again estimated to be 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 create 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, the bulk of the estimates were closer to zero than to one. It might seem, then, that these results rescue the Euler equation from the Scylla of a (rejected) prediction that $\alpha_1 = 0$ only to smash against the Charybdis of a (rejectable) prediction that $\alpha_1 = 1$. Fortunately, there is an escape hatch. The theoretical arguments and simulation evidence presented thus far do *not* necessarily imply a coefficient of 1 on $E_t[\Delta \log \tilde{Y}_{t+1}]$ – they imply a coefficient of one on $E_t[\Delta \log \tilde{P}_{t+1}]$. 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 panels 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’s *level* is temporarily low today, income *growth* between today and tomorrow is likely to be high, and vice versa. Hence the instrument used for $E_t[\Delta \log \tilde{Y}_{t+1}]$ is V_t .) Panels 5 and 6 reveal that the coefficient on predictable transitory movements in income is statistically significantly different from zero, but, at around 0.10, 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 predictable transitory and for predictable 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 decomposed that growth into the part representing transitory growth and the part representing permanent growth.¹⁵ Essentially all we can say is that (under this range of parameter values), the coefficient on predictable income growth should be somewhere between 0.10 and 1.0. Of the roughly twenty studies

¹⁵Note that the analysis here relies heavily on our assumption of an unchanging rate of growth for permanent labor income. If the growth process for P_t is more complicated than assumed here, it is not necessarily the case that the coefficient on predictable changes in P_t should be one.

cited by Browning and Lusardi (1996), none (to my knowledge) attempts to decompose predictable income growth into predictable transitory and predictable permanent components.¹⁶ Since the confidence intervals for α_1 in virtually all of these papers 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 of 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, they 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 an optimizing model like the one considered here 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), for example, 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, and 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

¹⁶Most of these papers use instruments such as occupation or education group as instruments. It might seem that such variables should be more highly correlated with permanent than with transitory income growth. However, there are well-documented differing patterns of cyclicity for different occupations and education groups. To the extent that the instruments capture such cyclical rather than secular movements, they will be predicting transitory as well as permanent growth. The ideal test would be to regress a long change in consumption on an instrumented lagged long difference in income, $\log C_{t+10}/C_t = \alpha_0 + \alpha_1 \log Y_{t+10}/Y_t$.

Euler equation, equation (8),

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

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 $\alpha_2 = \frac{1+\rho}{2}$.¹⁷ There is a widespread impression that, if any instruments can be found that are correlated with (η_{t+1}^2) , 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)

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

Assuming that 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 . In other words, the independent variation in η^2 caused by variation in r is perfectly correlated with the r_j whose coefficient is already being estimated.

One might hope that an instrument correlated with the impatience parameter δ_j could serve to identify α_2 . Certainly, an instrument correlated with impatience should generate variation in wealth and therefore in η_{t+1}^2 and so may look like a good instrument in first-stage instrument validity tests. And it is quite plausible to suppose that the time preference rate is correlated with observable variables such as, say, the consumer's level of education (one of the instruments typically used for η_{t+1}^2). The first panel of table 3 therefore presents the results when equation (16) is estimated on simulated 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 with an η_{t+1}^2 instrument (like education) that is correlated with the time preference rate is that δ_j also enters the Euler equation in another place: In the $\rho^{-1}(r_j - \delta_j)$ term. Econometrically, this means that δ_j is correlated with an unobserved variable ($\rho^{-1}\delta_j$) that is correlated with the included (instrumented) variable η_{t+1}^2 , a situation

¹⁷Dynan's regression actually includes lagged individual wealth as an independent variable, and therefore is not exactly equivalent to (16); see below for simulation results obtained when lagged wealth is included.

Panel	Consumers in Sample [†]	Instruments [‡]	Coe . on r_j		Coe . on $(\eta_{t+1}^2)_j$	
			Mean*	[.05-.95]	Mean*	[.05-.95]
1	All (10889)	RDUM+ DELDUM	-0.01 (0.15)	[-0.27,0.24] (0.14,0.17)	-0.85 (0.76)	[-2.17,0.31] (0.58,1.04)
2	BASE+R (4950)	RDUM+ GDUM	0.16 (0.36)	[-0.30,0.69] (0.19,0.66)	10.29 (5.28)	[5.14,18.34] (1.57,12.30)
3	All (7919)	All But RHODUM and DELDUM	-0.01 (0.14)	[-0.25,0.24] (0.13,0.14)	-0.42 (0.25)	[-0.79, -0.05] (0.23,0.27)

Notes: The dependent variable in all regressions is the growth rate of consumption, $\Delta \log C_{i,t+1}$.

[†]The column labelled 'Consumers In Sample' indicates which simulated consumers are included the sample. For example 'Base + R' in Panel 2 means that the sample includes consumers with baseline parameter values and those for whom R differs from the baseline.

[‡]The column labelled 'Instruments' indicates which categories of dummy variables are used as instruments for r and η_{t+1}^2 when the regression is estimated on the data generated by the simulations. For example RDUM indicates use of three dummy variables indicating which of the three possible interest rates the consumer faces.

*The first row in each panel presents the average value and range of the coefficient estimates across the Monte Carlo simulations. The second row in each panel presents the average value and range of values for the regressions' estimates of the standard error.

Table 3: Second-Order Approximation Estimated on Simulated Data

that implies that the coefficient estimate on η_{t+1}^2 will be biased. 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 the one Dynan performs) will not detect this problem because OID tests only find correlations of instruments with the dependent variable that are *not* captured by the variables that are included, but since $(\eta_{t+1}^2)_j$ is included the OID test should not reject the specification.

A simple thought experiment may clarify the problem better than the foregoing abstract 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; instrument validity test regressions of η_{t+1}^2 on the instruments will find that the instruments do 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 regression coefficient estimates on both r_j and $(\eta_{t+1}^2)_j$ will be zero.

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 completely uncorrelated with preferences. As a practical matter, it is likely to be hard to identify instruments for which a compelling case can be made that they are correlated with r_j and g_j but uncorrelated with preferences. However, there is of course no difficulty in

simulated data. The next panel of table 3 therefore 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, in panel 2, are interesting. While the coefficient on the interest rate term is still insignificant, the mean coefficient on the η_{t+1}^2 term is 10.3; since equation (8) implies that this coefficient is equal to $\frac{1+\rho}{2}$, this would indicate a coefficient of relative risk aversion of almost 20. (The mean standard error of 5.28 indicates that the typical regression estimate would be able to reject the ‘correct’ coefficient $(1 + \rho)/2 = 2$, at around the 10 percent level of confidence).

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, suggest 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 $E_t[\tilde{\eta}_{t+1}^2]$ 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 $E_t[\tilde{\eta}_{t+1}^2]$ 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 -0.42, and is no longer significantly different from zero – again reproducing Dynan’s result, as in panel 1.

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.¹⁸ Neither approach appears capable of identifying structural parameters even in a dataset consisting exclusively of consumers behaving exactly according to the model.

¹⁸Ludvigson and Paxson (2001) also find that IV estimation of the second-order equation like that employed here fails to identify the true parameter values in a population of simulated consumers. They typically find a downward bias in the coefficient on $(\eta_{t+1}^2)_j$, like the findings in panel 3, rather than an upward bias as found in panel 2.

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?

6.1 Bad Ideas

6.1.1 GMM Estimation

An obvious answer is to blame all of the foregoing pathologies on approximation error, implying that 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). The first panel of Table 4 presents the results of GMM estimation on the baseline set of simulated consumers.¹⁹ 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 panels of table 4. Panel 2 reflects the results when the same data on C_{t+1}/C_t that are used for panel 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. Panel 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 further.

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 panels of the table present the results that emerge if the econometrician falsely assumes that $\beta = 0.99$ (panel 4) or $\beta = 1/(1.08)$ (panel 5).²⁰ 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

¹⁹The table presents the results from a Monte Carlo analysis of 1000 GMM estimations on 1000 different random collections of consumers from the group with baseline preferences.

²⁰All that matters for these equations is the product $R\beta$, so separate experiments showing the results for incorrect assumptions about R would be redundant.

Panel	Problems [†]	GMM Estimate of ρ^{\ddagger}	[.05-.95] Range
1	None	3.06 (0.55)	[2.21,4.32] (0.39,0.81)
2	Signal/Noise = 1/2	1.77 (0.29)	[1.38,2.23] (0.24,0.36)
3	Signal/Noise = 1/3	1.18 (0.19)	[0.94,1.47] (0.16,0.23)
4	Assumed $\beta = 0.99$ when true $\beta = 1/1.04$	2.74 (0.58)	[1.84,4.03] (0.43,0.82)
5	Assumed $\beta = 1/1.08$ when true $\beta = 1/1.04$	3.42 (0.54)	[2.60,4.71] (0.37,0.81)

Notes: Direct estimation of ρ from $R \sum_i [(C_{i,t+1}/C_{i,t})^{-\rho}] - 1 = 0$.

[†]The column labelled 'Problems' indicates for each panel the nature of the empirical problem being explored with the simulated data. For example, 'Signal/Noise = 1/2' examines the effects of white noise measurement error with the same stochastic properties as the 'true' variation in consumption growth.

[‡]The first row in each panel presents the average value and range of the coefficient estimates across the Monte Carlo simulations. The second row in each panel presents the average value and range of values for the regressions' estimates of the standard deviation.

Results summarize 1000 Monte Carlo simulations.

Table 4: Euler Equation Estimated on Simulated Data Using GMM

Panel	Sample	r_j		g_j		Average NOBS
		Mean	[.05-.95] Range	Mean	[.05-.95] Range	
1	BASE + RDUM + GDUM	-0.02 (0.03)	[-0.06,0.03] (0.03,0.04)	0.07 (0.03)	[0.02,0.12] (0.02,0.05)	4950

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. Notation is similar to previous tables.

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

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.

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.²¹ The estimated coefficient on g is about 0.07 and is typically statistically significant; more than 95 percent of the time, the estimated coefficient on g is positive. But the point estimate implies value of $\rho = 2/.07 - 1 \approx 28$, compared with its true value of 3. The point estimate of the coefficient on the interest rate implies a value of about $\rho = 9.5$, but is statistically insignificant. Thus, equation (18) also fails to provide a consistent way to estimate ρ . Again, the problem is that higher-order terms in the Taylor expansion must be correlated with g_j and r_j , so that (18) is misspecified because there are omitted variables correlated with the included variables.

²¹The only consumers included were those from the baseline group and those groups for whom r or g varied from the baseline.

6.1.3 Individual-Specific Euler Equation Estimation

The arguments to this point in the paper have been directed at demonstrating that the Instrumental Variables approach to Euler equation estimation traditionally used in micro data 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.²² 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 relied on the proposition (verified by simulations) that the group mean values of the $\eta_{i,t+1}^2$ terms would take 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 $E_{i,t}[\tilde{\eta}_{i,t+1}^2]$, then it might be possible to estimate equation (8) *without* using instrumental variables. To be specific, one could estimate:

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

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 5.13, which implies an estimate of about $\rho = 9$ —an *upward* bias, like that found in the instrumental variables regressions reported earlier in panel 2 of table 3.

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, panels 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. Furthermore, as a final blow to the idea of determining the magnitude of the bias and adjusting for it, note that Laibson (1997) performs an experiment conceptually similar to the one examined here, and finds a *downward* bias to the estimate of ρ . I have been unable to determine why Laibson's results differ from those reported here (his model, parameters, and techniques differ somewhat, but which differences are crucial is unclear). But his differing results cast doubt on the possibility that it might be possible to get a robust and precise idea of the magnitude of the bias, and adjust for it—especially since in subsequent work, Laibson, Repetto, and Tobacman (1998) perform a similar experiment again, and find an upward bias!

As a way of investigating the source of the bias in my estimates, Figure 3 plots the true numerical expectation of $E_t[\Delta \log \tilde{C}_{t+1}]$ as a function of the level of cash-on-hand under

²²Except for the regressions examining the correlation between predictable transitory income growth and consumption growth, where idiosyncratic information on the level of the transitory shock in period t was used to predict income growth between t and $t + 1$.

Panel	Consumers in Sample [†]	Estimation Method [‡]	$E_{i,t}[\eta_{i,t+1}^2]$		Average NOBS
			Mean	[.05-.95] Range	
1	BASE	OLS	5.13 (0.99)	[3.56,6.80] (0.85,1.13)	990
2	R=1.00	OLS	4.49 (0.91)	[2.94,6.14] (0.76,1.04)	990
3	R=1.04	OLS	4.16 (1.03)	[2.52,5.78] (0.90,1.16)	990
4	$\beta = 1.00$	OLS	4.16 (1.23)	[2.23,6.07] (1.07,1.40)	990
5	$\beta = 1/1.08$	OLS	4.75 (0.85)	[3.40,6.15] (0.77,0.94)	990
6	$g = .02$	OLS	3.77 (1.69)	[0.88,6.54] (1.39,2.02)	990
7	$g = .06$	OLS	4.01 (0.68)	[2.92,5.15] (0.61,0.75)	990
8	$\rho = 1$	OLS	1.90 (0.29)	[1.42,2.36] (0.27,0.31)	990
9	$\rho = 5$	OLS	5.81 (1.21)	[3.78,8.10] (0.98,1.45)	990
10	$\sigma_v = .05$	OLS	2.64 (1.09)	[0.99,4.26] (0.91,1.31)	990
11	$\sigma_v = .15$	OLS	2.52 (0.40)	[1.77,3.26] (0.34,0.46)	990
12	$\sigma_n = .05$	OLS	3.12 (0.34)	[2.50,3.74] (0.29,0.38)	990
13	$\sigma_n = .15$	OLS	2.74 (1.46)	[0.45,5.08] (1.19,1.76)	990
14	BASE	IV	8.45 (4.03)	[4.92,14.36] (1.18,8.73)	990
15	R=1.00	IV	6.78 (2.53)	[4.04,11.10] (0.93,5.48)	990
16	R=1.04	IV	7.92 (6.36)	[3.73,14.90] (1.30,13.09)	990

Notes: The dependent variable in all regressions is the growth rate of consumption, $\Delta \log C_{i,t+1}$.

[†]The column labelled 'Consumers In Sample' indicates which simulated consumers are included the sample. For example 'R=1.00' in Panel 2 means that the sample includes only consumers for whom R=1.04 and all other parameter values are at their baseline values.

[‡]Panels 1-13 regress consumption growth for an individual consumer on that consumer's idiosyncratic expectation of $\eta_{i,t+1}^2$. Since this is an unobservable variable in real datasets, panels 14-16 present results when the ex-post value of $\eta_{i,t+1}$ is instrumented using $x_{i,t}$ and $x_{i,t}^2$.

Table 6: Second-Order Approximation Using Idiosyncratic Data

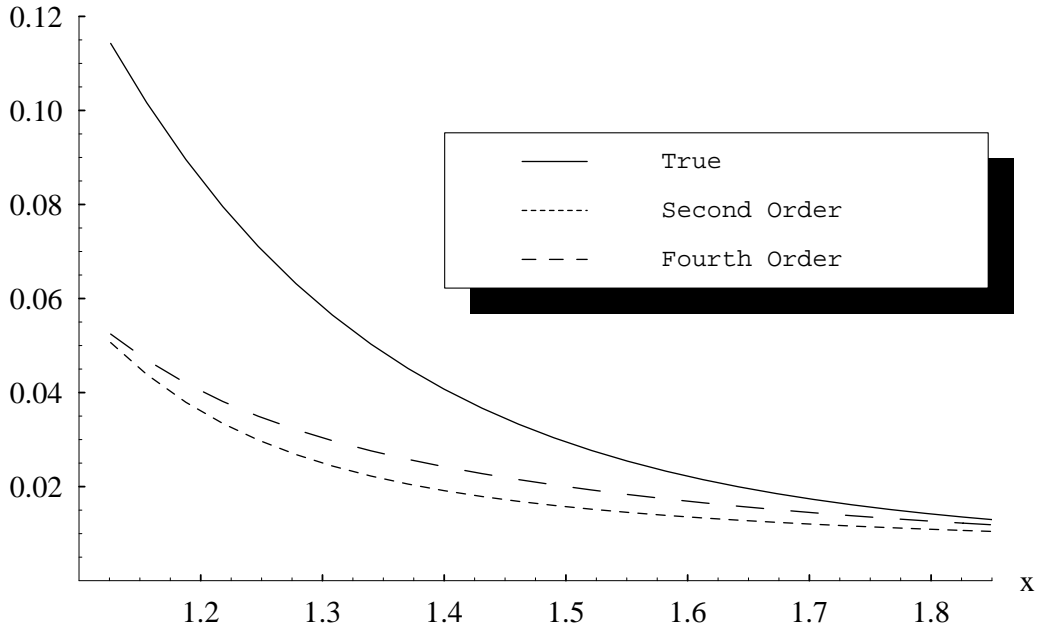


Figure 3: True and Approximated $E_t[\Delta \log \tilde{C}_{t+1}]$

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 $E_t[\tilde{\eta}_{t+1}^2]$ is not constrained to be $\frac{\rho+1}{2}$, the regression chooses a much larger value for that coefficient, with an `o` setting 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 actual data, it seems clear that it will not be possible to estimate the coefficient of relative risk aversion on a population of bu er-stock consumers with much precision using any plausible approximation to the consumption Euler equation.²³

²³It may seem puzzling that even a fourth order approximation performs so poorly; the problem lies in the poor quality of even a high-order Taylor approximation when it attempts to deliver the function's value at points very far from the point around which the approximation has been taken. Specifically, most of

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. Leonhard 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 near-impossibility of recovering a direct estimate of structural parameters from consumption growth regressions (at least on a cross-section microdata population with a significant proportion of bu er-stock consumers), 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 bu er-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 0.10 for baseline parameter values). These are eminently testable propositions.²⁴

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 $E_{i,t}[\tilde{\eta}_{i,t+1}^2]$ 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 highly statistically significant relationship between consumption growth and $E_{i,t}[\tilde{\eta}_{i,t+1}^2]$.

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, $E_{i,t}[\tilde{\eta}_{i,t+1}^2]$ 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 $E_{i,t}[\tilde{\eta}_{i,t+1}^2]$. Panel 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 $E_{i,t}[\tilde{\eta}_{i,t+1}^2]$ taken from the model.²⁵ Panels 15 and 16 show that similar results obtain for two of the other groups of

the error arises as a result of serious errors in approximating the value when the consumer is experiencing one of the rare but extremely potent zero-income events. For proof of this assertion, run the program `WhyIsTaylorExpnSoBad.m` included in the set of Mathematica files that reproduce the paper's results, available on my website.

²⁴A 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.

²⁵Ludvigson and Paxson (2001) also report that including the lagged level of wealth as an instrument increases the magnitude of the estimated coefficient on the $(\eta_{i,t+1}^2)_j$ term.

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 and data available.

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.²⁶ 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 (1994), 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 the one 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 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, several recent papers have scored impressive success in doing this. The pioneering work was by Parker and Gourinchas (1999) and Palumbo (1999); more recent contributions include work by Cagetti (1999) and French (2000). In all of these papers, the authors develop routines to quickly solve and simulate a dynamic

²⁶It 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	\bar{R}^2
		Mean	[.05-.95] Range		
1	R	2.33 (0.20)	[2.00,2.65] (0.20,0.21)	2971	0.04
2	G	-8.58 (0.24)	[-8.98, -8.18] (0.24,0.25)	2970	0.29
3	σ_n^2	16.70 (0.64)	[15.49,17.92] (0.60,0.67)	2970	0.19
4	σ_v^2	5.16 (0.48)	[4.33,6.00] (0.46,0.50)	2971	0.04
5	ρ	0.192 (0.002)	[0.188,0.195] (0.002,0.002)	2970	0.71

Notes: The dependent variable in all regressions is the level of cash-on-hand, $x_{i,t}$. The assumption is that actual values of all variables are directly observed. IV estimation is also possible and should produce consistent estimates. For each parameter examined, the sample consists of the base group plus the set of consumers whose value of that parameter is different from the baseline value.

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

life cycle 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 some empirical dataset.

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

7 Time Series Estimation of Euler Equations

Partly in response to the initial draft of this paper, Attanasio and Low (2000) have written a paper that argues that it is possible to use time series variation in interest rates estimate Euler equations successfully, either using panel data on individual households or using a time series of repeated cross sections. Their method is similar to that of this paper: Solve the model under certain assumptions about parameter values, simulate the behavior of a set of consumers behaving exactly according to the model, and estimate regressions on the resulting simulated data. They find that if the time series estimation covers a long enough span of time, they are able to obtain estimates of the IES that are fairly close to the ‘true’ IES used in solving the model.

Perhaps the most problematic aspect of their analysis is that they do not present any information on the sensitivity of their results to variation in any of several crucial parameters that one would expect to determine success or failure of the method.

One such parametric assumption has to do with serial correlation of the real interest rate. The importance of the serial correlation assumption can be understood by realizing that if interest rates never varied and consumers were impatient, consumption growth would equal income growth for all the reasons detailed above. Similarly, if interest rates tended to remain at the same level for a very long time, and then to suddenly move to a new level where they would again remain for a long time, it is clear that consumption growth would equal income growth almost all the time (except during the brief intervals during which the new ergodic distribution would be established after one of the rare interest rate shocks).

From the standpoint of maximizing the ability of the econometric method model estimate the IES, the ideal interest rate process is therefore one that has a lot of predictable variation. Thus, an AR(1) process in which the coefficient on lagged interest rates is somewhere well away from zero (because in that case there is no predictability to interest rates) and well away from one (because in that case there is no predictability to the change in interest rates) is what one would want.

Attanasio and Low (2000) use an AR(1) parameter of 0.6 for the real interest rate, which they indicate matches empirical data for the UK.²⁷ Even for this choice of serial correlation coefficient, Attanasio and Low find that it is necessary to have at least ten years worth of data ($T = 40$ in their notation) in order to obtain a reliable estimate of the intertemporal elasticity of substitution. For the US the serial correlation coefficient in real interest rates appears to be somewhat higher than the 0.6 figure reported by Attanasio and Low for the UK, and the US serial correlation coefficient also exhibits some evidence of varying over time. If their results are highly sensitive to the serial correlation coefficient, then the generality with which they can conclude that Euler equation estimation works using (enough) time-series data is compromised.

Attanasio and Low also make nonstandard assumptions in several other dimensions, and present little sensitivity analysis with respect to those assumptions. Their coefficient of relative risk aversion is fixed at $1/0.67 = 1.5$ throughout the paper, while most of the precautionary saving literature has tended to present results for parameter values centered on 3 and ranging up to 5. As demonstrated in several places above, the assumption about ρ can make a big difference to results (see, e.g., Table 6). And assuming a small precautionary motive obviously reduces the magnitude of the precautionary effects that interfere with Euler equation estimation.

Furthermore, there are several other dimensions along which the model examined by Attanasio and Low differs from the one presented here in ways that may make it easier to successfully estimate an IES. One of these is that Attanasio and Low assume that there is no income growth over the lifetime; empirical data show robust rates of income growth during most of the working life for most consumers. A rapidly growing level of income induces consumers to be ‘impatient’ in the sense required to generate buffer-stock saving behavior, so one might expect that the results that Attanasio and Low report for seemingly plausible values of the time preference rate might substantially understate the importance of the precautionary effects that arise as an interaction between impatience and prudence.

²⁷Although their text is a bit unclear, I will assume that this is at an annual rate; the points made below should be even stronger if their 0.6 represents a quarterly rate.

Indeed, when Attanasio and Low make consumers very impatient by assuming a high time preference rate, they find that the performance of their Euler equation estimation deteriorates (though they must assume very high time preference rates to make their estimates badly biased).

Also, Attanasio and Low do not allow for the possibility of transitory as well as permanent shocks to income. This may understate the short-term precautionary saving motive, and thereby reduce the magnitude of short-term deviations of the model with uncertainty from the certainty equivalent model.

It seems clear, therefore, that there will be ranges of parameter values where time-series estimation of the Euler equation will work and ranges where such methods will fail. The coming debate on time series estimation is therefore most fruitfully formulated as an argument about what the boundaries of those regions are, and whether plausible parametric configurations lie within or outside of those regions. With Wendy Dunn, I am currently working to explore this question.

8 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, at least if the estimation is performed on a population of consumers many of whom are engaged in buffer-stock saving behavior. 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.

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

All of the programs used to generate the results in this paper are available at the author's website, <http://www.econ.jhu.edu/People/ccarroll/carroll.html>.

References

- ATTANASIO, ORAZIO P., AND HAMISH LOW (2000): “Estimating Euler Equations,” *NBER Technical Working Paper Number 253*.
- ATTANASIO, ORAZIO P., AND GUGLIELMO WEBER (1995): “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey,” *Journal of Political Economy*, 103, 1121–57.
- BARSKY, ROBERT B., F. THOMAS JUSTER, MILES S. KIMBALL, AND MATTHEW 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,” *Journal of Economic Literature*, 34(4), 1797–855.
- CAGETTI, MARCO (1999): “Wealth Accumulation Over the Life Cycle and Precautionary Savings,” *Manuscript, University of Chicago*.
- CAMPBELL, JOHN Y., AND N. GREGORY MANKIW (1989): “Consumption, Income, and Interest Rates: Reinterpreting the Time-Series Evidence,” in *NBER Macroeconomics Annual, 1989*, ed. by Olivier J. Blanchard, and Stanley Fischer, pp. 185–216. MIT Press, Cambridge, MA.
- CARROLL, CHRISTOPHER 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), 111–148, <http://www.econ.jhu.edu/people/ccarroll/howdoesfuture.pdf>.
- (1996): “Buffer Stock Saving: Some Theory,” *Manuscript, Johns Hopkins University, January*, <http://www.econ.jhu.edu/people/ccarroll/BufferTheory.pdf>.
- (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, CXII(1), 1–56, <http://www.econ.jhu.edu/people/ccarroll/BSLCPIH.pdf>.
- CARROLL, CHRISTOPHER D., AND MILES S. KIMBALL (1996): “On the Concavity of the Consumption Function,” *Econometrica*, 64(4), 981–992, <http://www.econ.jhu.edu/people/ccarroll/concavity.pdf>.
- CARROLL, CHRISTOPHER D., AND ANDREW A. SAMWICK (1997): “The Nature of Precautionary Wealth,” *Journal of Monetary Economics*, 40(1), <http://www.econ.jhu.edu/people/ccarroll/nature.pdf>.

- CARROLL, CHRISTOPHER D., AND LAWRENCE H. SUMMERS (1991): "Consumption Growth Parallels Income Growth: Some New Evidence," in *National Saving and Economic Performance*, ed. by B. Douglas Bernheim, and John B. Shoven. Chicago University Press, Chicago.
- CARROLL, CHRISTOPHER D., AND DAVID N. WEIL (1994): "Saving and Growth: A Reinterpretation," *Carnegie-Rochester Conference Series on Public Policy*, 40, 133–192.
- DYNAN, KAREN E. (1993): "How Prudent Are Consumers?," *Journal of Political Economy*, 101(6), 1104–1113.
- FRENCH, ERIC (2000): "The Effects Of Health, Wealth, And Wages On Labor Supply And Retirement Behavior," *Working Papers Series, Federal Reserve Bank of Chicago, WP 2000-02*.
- GOURINCHAS, PIERRE-OLIVIER, AND JONATHAN PARKER (1999): "Consumption Over the Life Cycle," *Manuscript, Princeton University*.
- HALL, ROBERT E. (1978): "Stochastic Implications of the Life-Cycle/Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 96, 971–87.
- HANSEN, LARS M. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, pp. 1029–1054.
- LAIBSON, DAVID (1997): "Hyperbolic Discount Functions and Time Preference Heterogeneity," *Manuscript, Harvard University*.
- LAIBSON, DAVID, ANDREA REPETTO, AND JEREMY TOBACMAN (1998): "Self-Control and Saving For Retirement," *Brookings Papers on Economic Activity*, 1998(2), 91–196.
- LAWRANCE, EMILY C. (1991): "Poverty and the Rate of Time Preference: Evidence from Panel Data," *Journal of Political Economy*, 99(1), 54–77.
- LUDVIGSON, SYDNEY, AND CHRISTINA PAXSON (2001): "Approximation Bias in Linearized Euler Equations," *Review of Economics and Statistics, Forthcoming*.
- MICHAELIDES, ALEXANDER, AND SERENA 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, MICHAEL G. (1999): "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle," *Review of Economic Studies*, 66(2).
- RUNKLE, DAVID E. (1991): "Liquidity Constraints and the Permanent Income Hypothesis: Evidence from Panel Data," *Journal of Monetary Economics*, 27(1), 39–98.
- SHAPIRO, MATTHEW D. (1984): "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data," *Economics Letters*, 14(1), 93–100.

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