

TECHNICAL WORKING PAPER SERIES

APPROXIMATION BIAS IN LINEARIZED  
EULER EQUATIONS

Sydney Ludvigson  
Christina H. Paxson

Technical Working Paper 236  
<http://www.nber.org/papers/T0236>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
March 1999

We thank Orazio Attanasio, John Y. Campbell, Christopher Carroll, Angus Deaton, Bo Honore, David Laibson, Martin Lettau, Steve Zeldes, and workshop participants at the NBER Aggregate Implications of Microeconomic Consumption Behavior Meeting, the NY Federal Reserve, and Princeton University for helpful comments, and Alex and Leo Kontorovich for research assistance. Paxson received support from the National Science Foundation for this work. Any errors and omissions are the responsibility of the authors. Any opinions expressed are those of the author and not those of the Federal Reserve Bank of New York, the Federal Reserve System, or National Bureau of Economic Research.

© 1999 by Sydney Ludvigson and Christina H. Paxson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Approximation Bias in Linearized Euler Equations  
Sydney Ludvigson and Christina H. Paxson  
NBER Technical Working Paper No. 236  
March 1999

**ABSTRACT**

A wide range of empirical applications rely on linear approximations to dynamic Euler equations. Among the most notable of these is the large and growing literature on precautionary saving that examines how consumption growth and saving behavior are affected by uncertainty and prudence. Linear approximations to Euler equations imply a linear relationship between expected consumption growth and uncertainty in consumption growth, with a slope coefficient that is a function of the coefficient of relative prudence. This literature has produced puzzling results: Estimates of the coefficient of relative prudence (and the coefficient of relative risk aversion) from regressions of consumption growth on uncertainty in consumption growth imply estimates of prudence and risk aversion that are unrealistically low. Using numerical solutions to a fairly standard intertemporal optimization problem, our results show that the actual relationship between expected consumption growth and uncertainty in consumption growth differs substantially from the relationship implied by a linear approximation. We also present Monte Carlo evidence that shows that the instrumental variables methods commonly used to estimate the parameters correct some, but not all, of the approximation bias.

Sydney Ludvigson  
Federal Reserve Bank of New York  
Domestic Research Function - 3F  
33 Liberty Street  
New York, NY 10045  
sydney.ludvigson@ny.frb.org

Christina H. Paxson  
Department of Economics  
and Woodrow Wilson School  
2189 Bendheim Hall  
Princeton University  
Princeton, NJ 08544  
and NBER  
cpaxson@wws.princeton.edu

## I. Introduction

The estimation of key parameters in the household's utility function has been a longstanding goal of empirical research on consumer behavior. These parameters are of interest because they furnish a quantitative gauge of theoretically important concepts, such the strength of the precautionary saving motive, the degree of prudence and risk aversion, and the willingness of consumers to substitute consumption over time.

In order to estimate parameters from the utility function, a suitable empirical specification must first be chosen. A wide range of applications employ specifications that rely on linear approximations to dynamic Euler equations--a virtual cottage industry has been created from the estimation of linearized consumption Euler equations.<sup>1</sup> Despite the popularity of this approach, an important unanswered question is whether linear approximations of Euler equations produce accurate estimates of key parameters in the consumer's utility function.

In this paper, we investigate the properties of linearized Euler equations within the context of one type of empirical application: the large and growing literature on precautionary saving that examines how consumption growth and saving behavior are affected by uncertainty. This application typically involves regressing consumption growth on measures of uncertainty in expected consumption growth, the idea being that, with precautionary motives, future uncertainty will depress current consumption and raise consumption growth. The estimating equation is derived from a second-order Taylor expansion of the Euler equation which is the first order condition for optimal consumption choice, and relates marginal utility today to expected marginal utility tomorrow. The parameter estimates can be used to measure the strength of

---

<sup>1</sup>To cite just a few examples, see Hall [1988]; Campbell and Mankiw [1989]; Campbell and Deaton [1989]; Zeldes [1989]; Kuehlwein [1991]; Dynan [1993]; Attanasio and Weber [1993]; Lusardi [1996]; Merrigan and Normandin [1996]; Ludvigson [forthcoming].

precautionary saving motives, where in the absence of precautionary motives future uncertainty should not affect consumption growth.<sup>2</sup>

The empirical work on precautionary saving has produced some anomalous results. Whereas some studies which have investigated the effect of income risk on the level of consumption or wealth suggest that precautionary motives may explain a significant fraction of wealth accumulation (for example, Skinner [1988]; Lusardi [1993]; Carroll [1994]; Carroll and Samwick [1995]), the estimated effects of consumption uncertainty on consumption growth are typically small, indicating that precautionary motives are weak or nonexistent (for example, Dynan [1993]; Kuehlwein [1991]). For utility functions characterized by decreasing absolute risk aversion, these latter results also imply implausibly low levels of relative risk aversion. More specifically, given that the within-period utility function is isoelastic, such that  $u(C_t) = \frac{C_t^{1-\rho}}{1-\rho}$ , this literature yields estimates of  $\rho$  that are generally below 1.3, and are often insignificantly different from zero.

We investigate one possible reason for these small estimates of  $\rho$ . Specifically, the method of estimating  $\rho$ —which utilizes a second order Taylor expansion of the Euler equation—relies on linear or loglinear *approximations* of the Euler equation. If the Euler equation is sufficiently nonlinear, then these approximations will be poor, implying that estimates of  $\rho$  may differ from their true values simply because the Euler equation is approximated. We refer to this divergence between the estimated and true values as “approximation bias”.

---

<sup>2</sup>As we discuss below, the impetus for precautionary saving in this literature is the existence of “prudence” in preferences, first defined by Kimball (1990). In practice, other factors may also lead to precautionary accumulation, such as the possibility that borrowing constraints may bind in the future (see Deaton, 1992).

In this paper, we study how approximation bias due to linear approximation influences parameter estimates. We take a two step approach. First, we investigate how nonlinear are the Euler equations, an analysis that can be performed without simulating data. We start with a standard intertemporal optimization problem, in which a finite-lived consumer with isoelastic preferences chooses consumption and saving given current wealth and the current shock to income. For each year of “life,” and for a variety of assumptions about the parameters that govern preferences and the income process, we solve numerically for the function that relates consumption to wealth and the income state. We can then contrast the actual relationship between expected consumption growth and uncertainty in consumption growth, with the relationship implied by the second order approximation. This step allows us to document how much the approximate relationship differs from the true relationship across a range of wealth values, income states, and parametric assumptions. We find that this difference is quite large for some wealth-income states under some sets of parameters values.

How might the nonlinearities we find influence actual parameter estimates? The second step in our analysis investigates this issue. In particular, we ask how nonlinearities are likely to be translated into approximation bias in estimates of  $\rho$ . Typical estimates based on linearized Euler equations involve using household-level data to regress consumption growth over a given time period on a measure of variability in consumption growth, such as average squared consumption growth, using either ordinary least squares or instrumental variables techniques. In the latter case, squared consumption growth is often instrumented using variables such as education and occupation indicators; the assumption is that these factors predict variability in consumption growth but are uncorrelated with the error term. The second step in our analysis

allows us to investigate the extent to which this assumption is likely to be true.

We examine the extent of bias using both ordinary least squares (OLS) and instrumental variables (IV) estimators by performing a Monte Carlo study in which we choose a utility function (i.e. a value of  $\rho$ ), and simulate panel data on optimal consumption for a large number of “households.” For our results to be informative about the extent of bias in actual parameter estimates, it is important that the heterogeneity across our simulated households mimics that found in true data. To do this, we use income data from the *Panel Study of Income Dynamics* (PSID) to estimate the parameters of the income process for households in sixteen different education/occupation groups. We then solve for the appropriate consumption functions using each of the sixteen sets of parameter estimates. With these consumption functions in hand, it is straightforward to simulate data for “households” from the different education/occupation groups (with the fraction of households from each group in the simulated data equal to the fraction of households in each group in the PSID). These data are used to investigate the properties of OLS estimates of  $\rho$ , as well as IV estimates that use occupation and education indicators as instruments.

Our results indicate that regressions of consumption growth on consumption growth squared produce estimates of  $\rho$  that are biased down from the true value. OLS estimates produce the most bias, with estimates of  $\rho$  that are between 12% and 30% of the true value. The IV estimates fare somewhat better, but are still biased down, with estimates of  $\rho$  that are typically around 60% of the true value. The reason for the bias is that the instruments (i.e. the occupation and education indicators that determine the income process parameters) are correlated with the higher order moments of consumption growth that are in the error term of the linearized

equation.

Another feature of our results is that the extent of bias varies with wealth: when the sample is split between “low wealth” and “high wealth” households, the estimates of  $\rho$  are biased down most for the “low wealth” group, and for some estimation techniques are actually biased up for the “high wealth” group. Differences in estimates of  $\rho$  across wealth groups is often taken as evidence that poorer households are more liquidity constrained than wealthier households. Our results indicate that these differences may also arise as the result of approximation bias.

Section II discusses in more detail how approximations to Euler equations have been used in previous literature, and why estimates based on these approximations may be biased. Section III describes our methods for computing consumption functions and shows the results of these computations. In this section, we investigate the extent to which the Euler equation is nonlinear by documenting how the approximate relationship between consumption growth and the variance of consumption growth differs from the true relationship across wealth values and income states. Section IV describes our estimation of income processes for each occupation/education group using household data, and discusses the simulation results. Section V concludes.

## II. Approximations to Euler Equations

We start with a simple but general model of consumption. Individuals choose consumption and saving in each period so as to maximize expected lifetime utility. We assume that there is one asset,  $A_t$ , and that assets held between  $t$  and  $t+1$  earn a gross return of  $R_{t+1}$ . Decisions are made conditional on current resources (cash-on-hand) held at the beginning of the time period, and on

information about future incomes and interest rates. Utility is additively separable and is discounted across periods at rate  $\delta$ . Sub-utility functions in each period are identical and isoelastic-elastic. The maximization problem is summarized as:

$$\text{Max } E_t \sum_{j=0}^{T-t} \left(\frac{1}{1+\delta}\right)^j \left(\frac{1}{1-\rho}\right) C_{t+j}^{1-\rho} . \quad (1)$$

Cash-on-hand ( $x_t \equiv A_t + y_t$ ) evolves according to:

$$x_{t+1} = R_{t+1} [x_t - C_t] + y_{t+1}, \quad (2)$$

where  $y_{t+1}$  is labor income earned in time  $t+1$ .

Following most of the literature on precautionary saving, we assume that the real interest rate is not stochastic and is fixed at  $r = R - 1$ . The only uncertainty consumers face is in labor income, which fluctuates from period to period. In this case, the Euler equation associated with utility maximization is:

$$u'(C_t) = \left(\frac{1+r}{1+\delta}\right) E_t [u'(C_{t+1})], \quad (3)$$

or, in the specific case of isoelastic-elastic utility:

$$C_t^{-\rho} = \left(\frac{1+r}{1+\delta}\right) E_t [C_{t+1}^{-\rho}]. \quad (4)$$

Because marginal utility is not linear in consumption in equation (4), it is not possible to derive



an equation that relates expected future consumption to current consumption. Instead, researchers commonly linearize the right-hand-side of (4) and derive an equation that relates the expected growth in consumption to the expected squared growth in consumption. Specifically, taking a second-order Taylor approximation of marginal utility in  $t+1$  around the point  $C_t$ , inserting into (4), and rearranging yields:

$$E_t \left[ \frac{C_{t+1} - C_t}{C_t} \right] = \frac{1}{\rho} \left[ \frac{r - \delta}{1 + r} \right] + \left( \frac{1 + \rho}{2} \right) E_t \left[ \left( \frac{C_{t+1} - C_t}{C_t} \right)^2 \right] + v_t, \quad (5)$$

where  $1 + \rho$  equals the coefficient of relative prudence  $-\frac{u'''(C_t)C_t}{u''(C_t)}$  as defined by Kimball (1990).

The error term,  $v_t$ , is composed of an additive series of moments of consumption growth.<sup>3</sup>

Specifically:

$$v_t = \sum_{j=3}^{\infty} (-1)^j \left[ \prod_{m=1}^{j-1} (m + \rho) \left( \frac{1}{j!} \right) \right] E_t \left[ \left( \frac{C_{t+1} - C_t}{C_t} \right)^j \right]. \quad (6)$$

Equation (5) indicates that, if precautionary saving motives exist, then uncertainty (as measured by the conditional expectation of squared future consumption growth) is positively related to anticipated consumption growth. The intuition underlying this result is straightforward: prudent individuals will delay consumption until uncertainty about the future is

---

<sup>3</sup>An alternate approximation can be obtained by assuming consumption is lognormally distributed, implying that  $\Delta \ln C_{t+1} = \rho^{-1}(r - \delta) + 1/2\rho \text{VAR}_t(\Delta \ln C_{t+1})$ , where  $\text{VAR}_t$  is the time  $t$  variance. Under this distributional assumption, higher order moments do not enter the approximate loglinear Euler equation. Consistent with the findings we discuss next, however, normality tests on simulated consumption data (not reported) rejected the hypothesis that consumption growth is lognormally distributed when the driving process (income growth) is lognormally distributed.

resolved, so that consumers facing more uncertainty will display higher consumption growth on average.

Equation (5) has been used as the basis for a large and growing body of empirical work. One set of papers examines the relationship between consumption growth and *income* risk. As has been pointed out by Dynan (1993), equation (5) concerns the relationship between consumption growth and uncertainty in *consumption*, not in income. However, it makes sense to establish, as an empirical fact, whether or not those who face riskier income streams have higher average consumption growth than others. These papers generally find that future income uncertainty decreases the *level* of current consumption, lending some support to the hypothesis that precautionary motives reduce the willingness of individuals to consume out of uncertain future income.<sup>4</sup>

An alternative approach is to estimate (5) directly. This is done by Dynan (1993), who uses household-level consumption data from the 1985 *Consumer Expenditure Survey*. This survey has a short panel element, with each household surveyed in as many as 4 consecutive calendar quarters. Dynan estimates (5) by regressing average consumption growth over the period for each household on (time) average squared consumption growth. She uses instrumental variables to account for the fact that taste shifters, which are likely to be correlated with the variance of consumption growth, may be included in the error term, and because the short length

---

<sup>4</sup>For example, Carroll (1994) presents evidence using a normalized variance of individual income, and Kimball's (1990) "equivalent precautionary premium" as measures of income uncertainty. Guiso, Jappelli, and Terlizzese (1992) provide similar evidence using a self reported measure of earnings uncertainty drawn from the 1989 Italian Survey of Household Income and Wealth. Note that these studies do not depend on linearized Euler equations, but instead rely on some reduced form solution to the consumption function itself.

of the panel will make the sample mean of actual squared consumption growth a poor measure of risk. The instruments include variables such as indicators for occupation and industry and education measures, which are plausibly related to consumption uncertainty. Dynan's fairly precise estimates of the coefficient on squared consumption growth range from .012 to .156, implying a coefficient of relative prudence in the range of .024 to .312, and a *negative* value of  $\rho$  (which also equals the coefficient of relative risk aversion) in the range of -.976 and -.688. As Dynan points out, this range of values for  $\rho$  is implausible. Merrigan and Normandin (1996), using British data, and Kuehlwein (1991), using the US *Panel Study of Income Dynamics*, also estimate relatively low values for  $\rho$ , in the range of .78 to 1.33 for Merrigan and Normandin. Kuehlwein uses *OLS* estimation and obtains a higher value for  $\rho$ , equal to about 4.0, but it is not statistically different from zero.

If the approximation error,  $v_t$ , is correlated with uncertainty in consumption growth, regressions of (5) will produce biased parameter estimates if the equation is estimated using ordinary least squares. Moreover, instruments that are correlated with uncertainty in consumption growth may be correlated with third and higher order moments of consumption growth that make up the error term in (5). To build intuition, consider, for example, the first term in (6), which equals expected cubed consumption growth multiplied by  $-(1+\rho)(2+\rho)/6$ . If those who face more uncertainty in consumption growth also have consumption growth that is more skewed to the right, the first term in  $v_t$  and expected squared consumption growth will be negatively correlated. In this case, and ignoring higher-order moments in  $v_t$ , a regression of consumption growth on expected squared consumption growth will yield an estimate of  $(1+\rho)/2$  that is biased down. Although it is not possible to prove analytically that the bias will go in one

direction or the other, our numerical calculations discussed below support the idea that bias due to approximation error will typically yield estimates of  $\rho$  that are too low. It is important to emphasize that this bias is not a result of measurement error, data quality, or misspecification of preferences—even if all variables were accurately measured and the econometrician knew the individual’s true objective function, the bias would still exist—but is instead solely the result of approximation error.

### III. Numerical Solutions to Consumption Functions

To assess the extent to which approximation error is a problem, we numerically compute consumption functions, and contrast the “true” relationships between consumption growth, consumption uncertainty, and the expected real interest rate, with the relationships implied by the linear approximations. This is done for a variety of assumptions about the size of  $\rho$  and the stochastic process that governs the evolution of income.

We start with the model discussed above<sup>5</sup>. The first step is to choose a utility function and a stochastic process for income. We choose the commonly assumed isoelastic utility function, and assume that income growth follows a first-order moving-average process with the general form:

---

<sup>5</sup>In general, there is no closed form solution to the optimization problem presented in (1)-(2) with risky labor income. However, for specific utility functions, an analytical solution can be derived. The most notable example is the quadratic utility case. This is not a case we want to analyze in detail since linear marginal felicity functions preclude precautionary savings motives. As a robustness check however, we used our numerical approach to solve for the optimal consumption solution when utility is quadratic and the real rate of interest equals rate of time preference, and verified that it was equivalent to the analytical solution.

$$\ln(y_t) = \ln(y_{t-1}) + \mu + \epsilon_t - \phi \epsilon_{t-1}. \quad (7)$$

The choice of a first-order moving average process is roughly consistent with evidence from the micro data, see for example MaCurdy (1982), Abowd and Card (1989), and Pischke (1995).

Although these studies generally suggest individual income changes follow a MA(2) rather than MA(1) process, the latter is a good approximation which requires one fewer state variables to solve the model, greatly reducing computational complexity.

As shown in Deaton (1991), the ratio of consumption to income at time  $t$  will be stationary, and can be solved for as a function of two state variables, the ratio of cash-on-hand to income (denoted  $w_t$ ) and the income growth innovation,  $\epsilon_t$ .<sup>6</sup> Let  $\theta_t$  equal the ratio of consumption to income in  $t$  and  $z_t$  equal the ratio of income in  $t$  to last period's income, so that  $z_t$  equals

$e^{\mu + \epsilon_t - \phi \epsilon_{t-1}}$ . The Euler equation can be expressed as:

$$\theta_t(w_t, \epsilon_t)^{-\rho} - \beta \int \theta_{t+1}([1+r][w_t - \theta_t(w_t, \epsilon_t)]z_{t+1}^{-1} + 1, \epsilon_{t+1})^{-\rho} z_{t+1}^{-\rho} dF(\epsilon_{t+1}) = 0, \quad (8)$$

where  $\beta$  equals  $(1+r)/(1+\delta)$ .

---

<sup>6</sup>It is now well known that serial correlation in income growth complicates the procedure for solving for consumption functions, since it means that there will be two state variables (cash-on-hand and lagged income) rather than one (see Deaton [1991], [1992]). Nonstationarity in income further complicates matters, since it means that in practice quite a wide range of incomes (and, therefore, state variables) may be possible. To deal with these problems, we work with stationary ratios of variables, solving for the optimal level of consumption relative to income. This specification is computationally more convenient than solving for the level of consumption itself, because it reduces the range of possible values for cash-on-hand, and implies that the second state variable in the model is the *innovation* to income growth, rather than the lagged level of income.

The Euler equations are solved via backwards recursion, starting with a terminal time period  $T$  (which we set to 60.) We choose a grid of 500 values of  $w_T$ , and solve for a set of corresponding values of  $\theta_T$ . Although it will generally be the case that  $\theta$  is a function of the income innovation  $\varepsilon$  as well as the cash-on-hand to income ratio  $w$ , this is not true of the terminal period: assuming that the consumer dies with no net assets, then consumption in the last period of life, period  $T$ , equals cash-on-hand, so that  $\theta_T$  equals  $w_T$  for all values of  $\varepsilon_T$ . Given this solution for  $\theta_T$ , one can solve for the values of  $\theta_{T-1}$  that satisfy the Euler equation for a grid of possible values of  $w_{T-1}$  and  $\varepsilon_{T-1}$ . Solutions for earlier periods can be found by working backwards, i.e. solving for  $\theta_t$  as a function of  $w_t$  and  $\varepsilon_t$ , given the solutions for  $\theta_{t+1}$ . Rather than assuming that the error  $\varepsilon_t$  has a continuous distribution, we specify the distribution as a discrete 10-point approximation to a normal distribution, which considerably speeds computation time.<sup>7</sup>

Figure 1 shows results from what we call our “baseline case.” In this case, we have set the parameters of the income equation (7) to  $\mu$  equal to .02,  $\varphi$  equal to .444, and the standard deviation of  $\varepsilon$  (denoted  $\sigma$ ) to .20. The parameters values for  $\mu$  and  $\varphi$  are from MaCurdy (1982) and Pischke (1996). Although their estimate of  $\sigma$  is .25, we use a lower value of .20 since at least part of the estimated error variance is likely to be the result of the presence of substantial measurement error in recorded income. The baseline interest rate  $r$  is set to .03, the rate of time preference,  $\delta$ , is set to .05, and the parameter  $\rho$  of the isoelastic utility function is set to 3.

Figure 1 shows characteristics of the consumption function for period  $t=30$ . The top left-

---

<sup>7</sup>In principle, it is possible to solve for consumption functions for any number of years. In practice, this is sometimes difficult since the range of possible values of cash-on-hand to income (eventually) narrows as the solution moves backward in time. In all of our cases we solved for at least 45 years, i.e., from  $t=16$  to  $t=60$ .

hand panel graphs the ratio of consumption to income against the cash-on-hand to income.<sup>8</sup> Note that as many as ten functions could be graphed, one for each of the ten possible values of  $\epsilon_t$ . To avoid clutter, we graph only three. The line marked “k=1” denotes the graph for the lowest value of  $\epsilon$ , “k=5” is the fifth-lowest value, and “k=10” is the highest. As expected, consumption is an increasing function of cash-on-hand. Furthermore, high realizations of  $\epsilon$  result in lower values of the ratio of consumption to cash-on-hand. This does *not* imply that good shocks to income lower consumption, but indicates that good shocks raise cash-on-hand by more than they raise consumption.

The top right-hand panel graphs the relationship between expected consumption growth and the cash-on-hand to income ratio. Note that consumption growth  $(C_{t+1} - C_t)/C_t$  can also be expressed as  $(\theta_{t+1}z_{t+1} - \theta_t)/\theta_t$ . The graph shows consumption growth expected as of time  $t$ :

$$\frac{E[\theta_{t+1}z_{t+1}|w_t, \epsilon_t] - \theta_t(w_t, \epsilon_t)}{\theta_t(w_t, \epsilon_t)}, \quad (9)$$

which is easily computed for each of the possible values of  $w_t$  and  $\epsilon_t$ . In the same way, we can compute expected squared consumption growth, which is shown plotted against cash-on-hand to income in the lower left-hand panel of Figure 1.

These two graphs indicate that a consumer with low cash-on-hand (relative to income) will have both higher expected consumption growth, and higher expected squared consumption growth. These results are consistent with the findings of Carroll (1997a) and Carroll and Kimball

---

<sup>8</sup> Figure 1 is drawn only for the range of cash-on-hand to income from .5 to 3; although nothing prevents cash-on-hand from being negative, simulation results indicate that it usually is not, and the graph focusses attention on the values of cash-on-hand that are most relevant.

(1996): because poor consumers have a lesser ability to smooth shocks to income, the conditional variance of their consumption growth will be higher than that for wealthier households, and precautionary motives work to depress consumption and increase its rate of growth.

Another feature of Figure 1 is that for any given cash-on-hand to income ratio, consumption growth is higher for lower values of  $\varepsilon_t$ . This ordering of the consumption functions is due to the existence of transitory noise in individual income. Because a negative shock to income today implies higher income tomorrow, individuals increase the amount of consumption relative to income out of any given level of cash-on-hand, thereby smoothing out transitory changes in income by accumulating and de-accumulating assets. However, individuals also expect higher consumption tomorrow, so the expected rate of growth of consumption increases.

Given the optimal solution for the first and second conditional moments of consumption growth, we are now in a position to compare their relationship with that implied by the Taylor expansion (5). This comparison is made in the bottom right panel of Figure 1, which graphs the relationship between expected consumption growth and expected squared consumption growth.<sup>9</sup>

The consumption growth equation implied by the Taylor expansion (5), indicates the function graphed “should” have an intercept equal to  $\frac{1}{\rho} \left[ \frac{r-\delta}{1+r} \right]$  and a slope of  $(\rho+1)/2$  given the baseline parameters. This linearized equation is also shown. To simplify notation in what follows, we denote the slope of the linearized equation as  $(\rho+1)/2$ , and the slope of the true consumption growth equation as  $(\rho^*+1)/2$ , although it should be kept in mind that  $\rho^*$  is not a

---

<sup>9</sup>There are actually ten different functions relating these two variables, one for each value of  $\varepsilon$ , and three of these are graphed. However, they are sufficiently similar that they are not visually discernable on the graph, so that they appear to lie on a single line.



fixed number, and varies with the cash-on-hand to income ratio.

A comparison of the two lines indicates that the linearized and true equations are quite different. In particular, the figure shows that the slope of the true equation is lower than the slope of the linearized equation. Figure 2 reproduces the lower-right hand panel of Figure 1 for a large number of time periods, and provides evidence that this result is not a peculiarity of the year chosen. In addition, the relationships in Figure 2 are graphed over a wider range of cash-on-hand to income ratios than in Figure 1. When expected consumption growth squared is very small, cash-on-hand is very large, and difference between the true equation and the linearized equation is small. This is most apparent in the figures for years later in life (years  $t-57$  through  $t-24$  for example) because the asset range for those years is large, taking on both very small and very large values.<sup>10</sup>

Two features of Figure 2 are important. First, the divergence between the slopes of the linearized and true functions varies with wealth. Specifically, the divergence is larger for higher values of expected consumption growth and consumption growth squared; this implies that the divergence in slopes is greatest for less wealthy consumers (for whom expected consumption growth is highest). It is therefore not surprising to find that estimates of  $\rho$  differ across groups of individuals when the sample is split based on wealth or occupation, as in Dynan (1993) and Merrigan and Normandin (1996). These studies consider the possibility that liquidity constraints may bias estimates of  $\rho$  downward for less wealthy consumers, and, consistent with this

---

<sup>10</sup>Note that asset ranges are not fixed from year to year and shrink as one works backwards in time. This implies that the figure for year  $t-12$  has a much smaller asset range than in figure  $t-57$ , and there is less variation in both expected consumption growth and expected consumption growth squared than there is later in life (e.g. in year  $t-57$ ). Thus, the figures are not entirely comparable since the axes vary with the range of cash-on-hand.

hypothesis, they report lower values for less wealthy households. Note that these findings are also consistent with the direction of bias shown in Figure 2, indicating that even without liquidity constraints, poorer households can be expected to appear less prudent simply because the linearized Euler equations of less wealthy individuals will be subject to greater downward approximation bias. Our simulation results, discussed below, confirm this prediction by showing that estimates of  $\rho$  differ across groups of individuals when the sample is split by wealth.

A second feature of Figure 2 is that the divergence between the linearized and true consumption growth equations is smaller for time periods that are farther back from the terminal year  $T$ . However, there is no evidence that the differences between the linearized and true equations disappear.

How large are the differences between the slopes of the linearized and true equations? To provide summary evidence, we computed the slopes of the true equations, shown in Figure 2, and then constructed weighted averages of the true slope  $(\rho^*+1)/2$  to the linearized slope  $(\rho+1)/2$ . The weights are constructed to reflect the distribution of the cash-on-hand to income ratio for each year that is implied by the model. Specifically, we simulated income series for 1000 consumers (using the baseline parameters for the income process), and used the policy functions for consumption to calculate series for consumption and cash-on-hand for each consumer. The simulated data is used to estimate the joint density of  $(w, \varepsilon)$  for each year.<sup>11</sup> This density function is used to compute the weighted average of the ratio of the true to linearized slope.

---

<sup>11</sup> The density is computed nonparametrically, using a quartic kernel with a bandwidth of 1.

The results for the baseline case are shown in the first column of Table 1. The top panel shows the average value of the true slope,  $(\rho^*+1)/2$ , divided by  $(\rho+1)/2$  for a selection of years. Consistent with Figure 2, this ratio increases as one moves backwards from the end of life, implying that the linear approximation is better in earlier time periods. However, most of this change in the ratio takes place in the last 10 time periods: between  $t=18$  and  $t=48$  the ratio declines by only .035. The lower panel of Table 1 shows the values of  $\rho^*/\rho$  implied by the results in the top panel; the value of  $\rho$  indicated by the (average) slope of the true consumption growth equation is between 68% and 82% of its true value.

How sensitive are our results to changes in the baseline set of parameters? Figure 3 shows results that correspond to those in the bottom right panel of Figure 1 (with  $t=30$ ), but with changes from the baseline case. Table 1 provides summary statistics for each of the cases for a selection of years. The results of Figure 3 and Table 1 are summarized as follows. First, the linearized Euler equation is a better approximation for lower values of  $\sigma$ , the standard deviation of  $\varepsilon$ : when  $\sigma$  is reduced from .20 to .125, the ratio of the true slope to that implied by equation (5) rises relative to the baseline case. For example, for  $t=30$ , the ratio equals .858 for the baseline case and .947 for the case with  $\sigma$  equal to .125, implying an increase  $\rho^*/\rho$  from .811 to .930. An implication is that estimates of  $\rho$  based on the linearized Euler equation may indicate that people with riskier income streams are less prudent (and less risk averse) than those with less risky income streams, even if in fact there is no difference across the groups. Bias due to approximation error may make it difficult to test whether risk aversion affects how individuals select into different occupations.

Second, reductions in  $\varphi$ , the moving average parameter, increase the divergence between

the slopes of the linearized and true equations. When  $\phi$  is lowered from .44 to .30, the ratio of the slopes (in  $t=30$ ) falls from .858 in the baseline to .793. Decreases in  $\phi$  imply that income innovations are more persistent, so that a given shock translates into a larger change in life-time wealth. It is therefore not surprising that declines in  $\phi$  have effects similar to increases in  $\sigma$ .

The third column of Table 1 shows the effect of increasing the parameter  $\rho$ . When  $\rho$  is increased from 3 to 4 the amount of bias rises. For example, the average ratio of the true slope to the linearized slope falls from .858 (for  $t=30$ ) to .770. Note this implies the value of  $\rho$  which would be inferred from the true equation is only 72.5 percent of the true value. Thus, the extent to which the degree of prudence is understated (in percentage terms) is positively related to the true degree of prudence.

The final two columns show the effects of changing the parameters  $r$  and  $\mu$ . The results indicate that increasing the rate of interest (and keeping the discount rate constant, so that consumers are more “patient”) has little effect on the extent of bias in the slope. Likewise, changing  $\mu$  has little effect.

The results in this section show that Euler equations are quite nonlinear, implying that, for some wealth values and income states, the true relationship between consumption growth and uncertainty in consumption growth will differ considerably from the linearized relationship.

#### IV. Approximation Bias in Empirical Estimates

This section examines whether errors caused by linear approximation to Euler equations result in biased parameter estimates of  $\rho$ , using commonly employed estimation techniques. The standard method used in the literature discussed in Section II is to construct measures of

consumption growth for household  $I$  in year  $t$  (denoted  $cg_{it}$ ) and measures of variability in consumption growth (denoted  $cg_{it}^2$ ) using panel data on households, and then regress consumption growth on variability in consumption growth:

$$cg_{it} = \beta_0 + \beta_1 cg_{it}^2 + \epsilon_{it}. \quad (10)$$

As discussed above, the consumption growth equation based on the linearized Euler equation implies that  $\beta_0$  equals  $\frac{(r-\delta)}{(1+r)\rho}$ , and that  $\beta_1$  equals  $(\rho+1)/2$ . Equation (10) is typically estimated using OLS or IV techniques. The use of instrumental variables makes a great deal of sense, if for no other reason than to eliminate biases caused by measurement error in consumption. However, the discussion in Section II indicates that it may be extremely difficult to find instruments that are not correlated with the error term in (10), even in the absence of measurement error.

Equation (6) indicates that the error term includes higher order moments of consumption growth. Valid instruments must consist of variables that are correlated with variability in consumption growth, but are not correlated with these higher-order moments. It is difficult to think of theoretically justifiable reasons why some variables would be correlated with the second moment but not with higher-order moments. Nevertheless, the extent of bias found in practice will depend on the extent to which the instruments are correlated with the error term, a factor that will, in turn, depend on which instruments are used.

The most common instruments employed in practice are variables such as occupation, education, and industry indicators (e.g., Dynan [1993]). The use of these variables is motivated by the idea that education and occupation influence the time-series properties of individual's income, so that (for example) more highly educated people may have higher income growth over their lifetimes, and self-employed managers may face more income risk than others. These

differences in income patterns will affect consumption decisions, and are likely to produce differences in variability in consumption growth. However, if (for example) individuals in different occupations and educations have different amounts of skewness in consumption growth, then these variables are not valid instruments, and even IV estimates of the parameters will be biased.

We assess whether this problem is important in practice by conducting a Monte Carlo analysis. Our approach is to generate simulated data on consumption for "households" who share common preferences and face the same interest rate, but have different values of the parameters for the process that generates income. In order to match heterogeneity in income parameters to heterogeneity found in real data, we use information on income from the *Panel Study of Income Dynamics* to estimate income equations for individuals in 16 different education/occupation groups.<sup>12</sup> We solve for up to 40 years of consumption functions using each of the 16 sets of parameters, and then simulate consumption data for households in each of the groups. The simulated data is used to estimate equation (10), above, using both OLS and IV techniques.

The first step in our analysis was to estimate the parameters of the income equation (7) for the different education and occupation groups. We used the 1972-1992 waves of the PSID, and selected a sample of male heads of households between the ages of 25 and 55, inclusive, who reported labor income for at least 3 contiguous years. Further details on sample selection are in Appendix A.

---

<sup>12</sup>In principle, it is possible to use information on industry, in addition to occupation and education, for each individual in the PSID. Placing households into categories that are as fine as those like occupation/education/industry resulted in cell sizes that are too small for empirical work. Thus, we chose to focus on groups defined by occupation/education which have larger cell sizes.

In each year, individuals were coded as being in one of the six occupations listed in Table 2. We then constructed a single non-time-varying occupation for each individual, that consisted of the occupation worked in the majority of years. (Ties were broken by assigning individuals to the "higher" of the occupations using the order in Table 2; for example, a person who spent 3 years as a manager and 3 years as a professional would be coded as a professional.) Individuals were also assigned to one of three education categories: less than 12 years of school, exactly 12 years of school, and more than 12 years of school. Individuals from two occupation/education cells were excluded from the analysis, since there were too few people.<sup>13</sup> The final data set consisted of 43,067 observations on 5,567 individuals. The distribution of individuals across occupation/education cells is shown in Table 2. Cell sizes ranged from 78, for professionals with exactly 12 years of education, to 719, for craftsmen with 12 years of education.

These data were used to estimate the parameters  $\mu$ ,  $\phi$ , and  $\sigma$  for each of the occupation/education groups. We followed the estimation methods discussed in Abowd and Card (1989); details are provided in Appendix A. The estimates for each cell are shown in Table 2. The estimates of  $\mu$  are what might be expected: within each occupation group, those with more education have higher earnings growth. Within each education group, craftsmen and operatives had the lowest growth. Somewhat surprisingly, laborers and service workers have higher values of  $\mu$  than do craftsmen and operatives. The parameter  $\phi$  also varies across groups; the general pattern is for those with more education to have lower values of  $\phi$ , indicating more persistence in shocks. Less skilled clerical and sales workers, and laborers and service workers, display

---

<sup>13</sup>These groups were professionals with less than 12 years of education, and managers with less than 12 years of education.

much less persistence. The parameter estimates of  $\sigma$  seemed too high to be plausible, and we halved them for the purposes of the simulations.<sup>14</sup> The halved values are reported in Table 2. There is remarkably little variation in  $\sigma$  across the cells. The range is from .137, for clerical and sales workers with less than 12 years of school, to .220, for laborers and service workers with more than 12 years of education. Values of  $\sigma$  cluster between .16 and .19 for nine of the sixteen cells. There is no clear tendency for  $\sigma$  to increase or decrease with education.

These parameters were used to solve for consumption functions for each of the cells. We assumed that the remaining parameters were constant across groups (with  $\rho=3$ ,  $r=.03$ , and  $\delta=.05$ .) and solved for 40 years of consumption functions. We then simulated data on income and consumption for 300,000 "households" for these 40 years, assuming that each individual started with cash-on-hand equal to income. The numbers of households in each occupation/education cell was set so that the distribution of households across cells in the simulated data matched that in the PSID. We discarded the first five years of data, because we did not want our results to be influenced by our assumptions about initial assets, and the last five years, when individuals are drawing down assets in anticipation of "death." The number of years per household was reduced by one more year after first-differencing consumption, leaving 8,700,000 household/year observations in our simulated data set.

Differences in the income parameters across occupation/education groups produce

---

<sup>14</sup>Several authors have pointed out that these measures are likely to be inflated by substantial measurement error in recorded income (for example, Deaton [1991]; Attanasio et al, [Forthcoming]). We also follow these authors in not reducing the size of the moving average parameter, even though this would logically follow in order to increase the role of the permanent component in income; the variance of the permanent component is already a large fraction of overall earnings variance, and a larger value seems implausible.



differences in consumption patterns in the simulated data. Table 2 shows averages of consumption growth and consumption growth squared for each cell of our simulated data. Average consumption growth ranges from 0.0016 (for clerical and sales workers with 12 years of education) to 0.796 (for highly educated professionals). These two groups also had the lowest and highest average risk in consumption growth.

To estimate the parameters of (10), we use Monte Carlo simulations to create 300 samples of 1000 households per sample, and used these samples to generate 300 sets of parameter estimates. Two versions of (10) were estimated. "Model A" refers to a regression of consumption growth between year  $t-1$  and  $t$  on squared consumption growth between the same period. The data are kept in panel form, with multiple observations per household. In "Model B," consumption growth and consumption growth squared were averaged over time for each household. This latter model is more similar to that estimated by Dynan (1993), who averages consumption growth and squared consumption growth over 4 quarters of the Consumer Expenditure Survey.

Both Model A and Model B were estimated using OLS, and using IV with three instrument sets. The first instrument set includes a set of dummies for each of the 16 occupation/education groups. The second set includes 6 occupation and 3 education dummies.<sup>15</sup> The third set is the same as the second, only it adds the lagged value of the ratio of assets to income. For model A, this is the value of the asset to income ratio at  $t-1$ ; for Model B, it is the value in the first year of the sample. The rationale for including this as an instrument is that

---

<sup>15</sup>It is common in empirical work to exclude interaction terms between the indicators in the instrument set.

consumers with more liquid assets may be better able to smooth consumption, and will have lower variability in consumption growth.<sup>16</sup>

As a final exercise, we split the samples for each model into two equally sized groups of consumers: those with "high" and "low" values of lagged ratio of income to assets. As discussed above, lower estimates of  $\rho$  for households with lower assets are taken as evidence of liquidity constraints. However, it is likely that approximation biases in the estimates also varies with wealth, and it is useful to see if the estimates of  $\rho$  vary across groups when, by construction, all households have the same value of  $\rho$  and no liquidity constraints exist.

Our primary interest is whether the slope parameter  $\beta_1$  differs from the values implied by the linearized equation, i.e.  $\beta_1=(\rho+1)/2=2$ . The OLS results, summarized in the top panel of Table 3, indicate that the estimated values of  $\beta_1$  are in most cases biased down. Model A, using the full sample, yields estimates of  $\beta_1$  with an average of .968 and a standard deviation of .389. The top left panel of Figure 4 shows a histogram for the 300 OLS estimates of  $\beta_1$ , and indicates that although there are many cases in which  $\beta_1$  is lower than .5, in no case does it exceed 1.7. The estimate of  $\rho$  implied by the average estimate of  $\beta_1$  is .936, less than one third of the true value of 3. The results for Model B (full sample) indicate even more bias. The average estimate of  $\beta_1$  equal to .688 implies an estimate of  $\rho$  of .376, only 12.5% of the true value.

The OLS estimates also perform badly when the sample is split by wealth. For Model A, the average estimate of  $\beta_1$  for the low-wealth sample is .864, lower than for the full sample. (The histogram of estimates for this case is shown in Figure 5.) Estimates for the high-wealth sample

---

<sup>16</sup>Dynan (1993) included the amount of interest and dividend income earned by the household during the 12 months prior to the survey as a proxy for liquid asset holdings.

are actually biased up: they all exceed 2, and imply on average a value of  $\rho$  equal to 3.9. The results for Model B are also yield estimates of  $\rho$  for the low-wealth sample that are lower than for the high-wealth sample, although for both groups there is downward bias.

Do IV estimators perform better than OLS estimators? Results using the three different instrument sets are in the second through fourth panels of Table 3, with histograms for each column of the table in Figures 4 through 9. The results can be easily summarized: First, the three sets of instruments produce very similar results. All of the instrument sets are good predictors of variability in consumption growth. Furthermore, within each model and sample, the instrument set makes little difference to the parameter estimates. Second, the IV estimates of  $\beta_1$  are generally larger than the OLS estimates. However, they are still biased down. For example, in Model A, the average IV estimate of  $\beta_1$  (using the first instrument set) is 1.408, in contrast to the average OLS estimate of .968. The IV results imply an estimate of  $\rho$  equal to 1.81, which is 60% of the true value. (Note that because the variables in the first two instrument sets do not vary over time, the full sample IV estimates for Model A and Model B are identical.) Third, the use of instrumental variables reduces, but does not eliminate, differences in the estimates of  $\beta_1$  across the wealth groups. Using the first instrument set, the values of  $\rho$  implied by the results are 1.77 for the low wealth group, and 2.10 for the high wealth group.

A common method of cross-checking the plausibility of a set of instruments is to test any overidentifying restrictions that have been made (for example, Dynan [1993]). We computed chi-square tests of the null that the overidentifying restrictions are valid. The number of rejections of the null were not frequent: they ranged from only 1 out of 300, to 45 out of 300 tests. We also examined whether the instruments are correlated with the third (noncentral)

moment in consumption growth. As discussed above, linearization of the Euler equation implies that cubed consumption growth is one of the components of the error term in (10). When using the full samples for Models A and B, the hypothesis that the instruments explained none of the variation in cubed consumption growth was rejected in at least 235 out of 300 cases. In particular, the estimation output of Model A (where there is a time series element to the sample) using instrument set 3 (which includes lagged cash-on-hand) shows that, although lagged cash-on-hand is a good predictor of expected consumption growth squared, it is also a good predictor of expected cubed consumption growth, a term that is included in the regression error of the linearized equation. Consequently, that specification produces rejections of the hypothesis that the instruments do not explain the variation in cubed consumption growth almost 100 percent of the time, yet the OI tests are rejected only 144 times out of 300 in the full sample. The results demonstrate that an inability to reject over identifying restrictions does not guarantee the validity of a set of instruments.

In summary, Monte Carlo analysis demonstrates that OLS estimation of consumption growth on uncertainty in consumption growth is likely to produce estimates of  $\rho$  that are biased down, and that IV estimation may eliminate some, but not all, of this bias. Recently, there has been a growing interest in investigating the extent to which linearizing consumption Euler equations might influence researcher's ability to recover parameters, and it is useful to compare our results to those of others.

Carroll (1997b) addresses this issue by simulating a model of individual consumer behavior, solved for three alternative values of the interest rate, the rate of time preference, the mean rate of permanent income growth, and the variances of permanent and transitory

components of income growth. He creates household groups of equal population which vary according to these parameter values, and then performs instrumental variables regressions of consumption growth on consumption growth squared using dummy variables that indicate group membership as instruments. In general, he finds a far greater degree of approximation bias in  $\rho$  than we find, and both the magnitude and direction of bias he finds is highly dependent on the empirical specification he analyzes. This may be due in part to the fact that he considers a different model of consumer behavior than that which we investigate, one often referred to as a "buffer stock" model. It differs from the present framework by imposing an impatience requirement: households prefer to do most of their consuming early in life.<sup>17</sup>

Other researchers have investigated the extent of bias in  $\rho$  in models that also differ from ours, but find results that are closer to those we present. Laibson (1997) investigates a standard buffer stock consumption model of the type considered by Carroll (1997a) and Deaton (1991) and finds that the value of  $\rho$  implied by the linearized regression is about 80 percent of its true value. Laibson, Repetto and Tobacman (1998) find that  $\rho$  is biased up by about 7 percent. While these results are closer in magnitude to our own, it is not surprising that they vary to some degree from our findings because there are meaningful differences between the model we consider and those that Laibson (1997) and Laibson, Repetto and Tobacman (1998) investigate; both of those studies investigate buffer stock behavior and so impose the impatience restriction. In addition, those studies assume an first order autoregressive process for the level of income, rather than an

---

<sup>17</sup>This requirement places a restriction on the values of key parameters--such as the standard deviation of income growth, the rate of time preference relative to the real interest rate, and the coefficient of relative prudence. This restriction makes the buffer stock model considerably less general than the model considered here.

MA(1) for the first difference of income.<sup>18</sup>

## V. Conclusion

This paper investigates nonlinearities in dynamic Euler equations. We also study the extent to which nonlinearities may induce bias in parameter estimates obtained when the estimating equation is a linear approximation to the nonlinear equation. We find that, for some wealth and income states, the Euler equation is sufficiently nonlinear that linearized approximations are poor and may imply a value for the coefficient of relative prudence (or risk aversion) that is lower than the true value. Simulation results confirm the prediction that ordinary least squares regressions of consumption growth on uncertainty in consumption growth are likely to produce a downward bias in the estimate of  $\rho$ . Moreover, the degree of bias varies with wealth, with less wealthy households displaying more downward bias. In addition, our results indicate that researchers may be able to eliminate some, but not all, of the bias that would arise in ordinary least squares estimation by instrumenting for consumption risk using education/occupation

---

<sup>18</sup>The differences in results between ours and Carroll's may also be attributable to the different estimators we consider. We have argued above that bias in the actual estimators will be a function of both the estimation technique and the particular choice of instruments, implying that it is important that heterogeneity across simulated households mimic that found in true data. Indeed, Carroll's results are quite sensitive to which set of group dummies he uses as instruments. But because his instruments are not observed and are not benchmarked on the data, it is difficult to know how the biases he finds may be translated into biases in practice. By contrast, the Laibson, Repetto and Tobacman study follows an approach more similar to ours by creating group dummies that are benchmarked to data and used in actual empirical studies: they create three education groups by estimating their income process using the PSID over education groups. This may explain why our results are in the same ball park as those of Laibson, and Laibson, Repetto, and Tobacman despite the fact that we investigate different models of consumer behavior.

dummies.

Our findings help to resolve some puzzles in literature on consumption behavior. For example, much of the empirical literature on precautionary saving suggests that income risk is an important factor in determining how much wealth consumers hold or how much they save, indicating that precautionary motives may be important. Yet studies which use linearized Euler equations often find that uncertainty in consumption growth has very little impact on consumption growth, suggesting that precautionary motives may be weak. Our results indicate that approximation error is likely to result in measures of the degree of prudence that are biased down. The analysis may also help explain why estimates of parameters of the utility function (using micro-level data) often differ across sub-samples of the population, split according to wealth or the degree of income uncertainty. Bias associated with the use of linear approximation varies with wealth and income risk, so that what would appear to be genuine differences in behavior across sub-groups may be an artifact of approximation error.

## Appendix A: Estimation of income parameters.

### 1. Data definitions and sample selection

The sample was drawn from the 1972-1992 waves of the *Panel Study of Income Dynamics*. For an observation to be included in the sample, an individual had to be a male head of household between the ages of 25 and 55, inclusive, in the survey year. The individual also had to have positive labor income, annual work hours that were positive but no greater than 4,680, have non-missing occupation and education information, and have average hourly labor earnings (i.e. annual labor earnings/annual hours) that were greater than or equal to \$1 and less than or equal to \$100. (Money values are in 1985 dollars.) After these exclusions were made, individuals were kept if they were in the sample for at least three contiguous years. The final sample had 43,067 observations on 5,567 individuals

The income measure used for the analysis was total labor income. Labor income is a variable constructed by the PSID, and equals includes income from wages and salaries, bonuses, overtime, commissions, a professional trade or practice, the labor components of farm and business income, and income from roomers and boarders and market gardens.

The occupation measure is based on the question on “occupation on the main job.” Occupation was given a 1-digit code in the early years of the survey (i.e. pre-1976), and the coding was then altered to a 2-digit and then to a 3-digit code. The occupation measures from the later years were recoded to be consistent with the earlier 1-digit measure; the possible occupation categories are in Table 2. We then defined our measure of occupation to be the “major occupation,” i.e. the occupation in which the person worked the most number of years. Ties were broken by allocating people to the “higher” of the occupations for which there was a tie, where “higher” occupations are listed earlier in Table 2. There is a moderate amount of occupational change: of the 5,567 people represented in the sample, 37.93 percent never changed occupation (so that major occupation always equaled current occupation). However, only 27% had more than 2 changes in occupation, and for many of these the changes were back and forth between only 2 occupations. Within each occupation group, we classified people according to whether they had less than 12 years of schooling, 12 years of schooling (i.e. a high school degree), and more than 12 years of schooling. The sample size did not permit classifying people into finer educational categories.

### 2. Estimation methods

The model we estimated for the individuals in each occupation/education group was:

$$v_{it} = \Delta \ln(y_{it}) - \mu = \epsilon_{it} - \phi \epsilon_{it-1} \quad (\text{A1})$$



where it is assumed that the shocks  $\epsilon_{it}$  are not correlated across time for individuals, have a zero mean, and a constant variance of  $\sigma^2$ . Equation (A1), together with these assumptions, implies that:

$$\text{var}(v_{it}) = \sigma^2(1+\phi) \quad (\text{A2})$$

and

$$\text{cov}(v_{it}, v_{it-1}) = -\phi\sigma^2. \quad (\text{A3})$$

We used the following strategy to estimate the parameters  $\mu$ ,  $\sigma$ , and  $\phi$ . First, following Abowd and Card (1989), who worked with experience-adjusted wage measures, we first regressed the change in the logarithm of labor income on a set of age dummies to sweep out age effects. (Since all individuals in the group are in the same education category, the age effects are close to being the same as experience effects.) We measured the growth term  $\mu$  as the intercept plus the age effect for 30-year-olds, i.e. the average change in the logarithm of income at age 30. An alternative would have been to measure  $\mu$  as the mean value of  $\Delta \ln(y_{it})$ . Doing so would have produced measures of  $\mu$  about half the value of the measures we used, but with the same ranking of growth across different occupation/education groups.

Second, we used the residuals from the regression described above to construct estimates of  $\text{var}(v_{it})$  and  $\text{cov}(v_{it}, v_{it-1})$ ; these estimates are simply the averaged squared residual, and the average of the residual times its lagged value. These were substituted into (A2) and (A3), respectively, and we then solved for  $\sigma^2$  and  $\phi$ . The standard errors of these estimates were obtained using the delta method. As discussed in the text, we used values of  $\sigma$  that are half of those estimated, on the grounds that the estimates seem implausibly high, and are likely to reflect measurement error. Measurement error in income that is *i.i.d.* should produce estimates of  $\phi$  that are also too high, i.e. income innovations will appear to be less persistent than they actually are. We experimented with estimating  $\sigma^2$  and  $\phi$  under the assumption that income has an *i.i.d.* measurement error component with a variance that is 1/4 the value of  $\sigma^2$ . This resulted in plausible estimates of  $\sigma^2$ , but values of  $\phi$  that were extremely low and, in one case (for professionals with 12 years of school) negative. We decided that the most sensible strategy was to use our original estimates, halving the estimate of  $\sigma$ .

The estimates and standard errors of the parameter estimated are listed in the following Table, with standard errors in parentheses:

	$\varphi$	$\sigma^2$
Professional, technical and kindred		
12 years of education	0.1589 (0.523)	0.0940 (0.275)
more than 12 years of education	0.2878 (0.144)	0.0870 (0.031)
Managers (self-employed and employees)		
12 years of education	0.3416 (0.146)	0.1244 (0.038)
more than 12 years of education	0.3968 (0.090)	0.1083 (0.017)
Clerical and sales		
less than 12 years of education	0.4068 (0.518)	0.0747 (0.056)
12 years of education	0.8093 (0.280)	0.0895 (0.019)
more than 12 years of education	0.3139 (0.214)	0.1427 (0.069)
Craftsmen and kindred		
less than 12 years of education	0.3291 (0.165)	0.1541 (0.056)
12 years of education	0.3412 (0.118)	0.1157 (0.029)
more than 12 years of education	0.2737 (0.167)	0.1164 (0.055)
Operatives (transport and non-transport)		
less than 12 years of education	0.4317 (0.142)	0.1169 (0.024)
12 years of education	0.3661 (0.111)	0.1157 (0.025)
more than 12 years of education	0.2991 (0.218)	0.1259 (0.070)
Laborers and service workers		
less than 12 years of education	0.5626 (0.149)	0.1317 (0.022)
12 years of education	0.4642 (0.162)	0.1750 (0.046)
more than 12 years of education	0.2633 (0.247)	0.2066 (0.163)

## References

- Abowd, John, and David Card, 1989, "On the Covariance Structure of Earnings and Hours Changes", *Econometrica*, 57, 411-45.
- Attanasio, Orazio P., and Guglielmo Weber, 1993, "Consumption growth, the interest rate and aggregation", *Review of Economic Studies*, 60, 631-49.
- Attanasio, Orazio P, James Banks, Costas Meghir and Guglielmo Weber. Forthcoming. "Humps and Bumps in Life Time Consumption." *Journal of Business and Economic Statistics*.
- Campbell, John Y., and N. Gregory Mankiw, 1989, "Consumption, income and interest rates: reinterpreting the time series evidence", *NBER Macroeconomics Annual, 1989*, 185-216.
- Campbell, John Y., and Angus Deaton. 1989. "Why is Consumption so Smooth?" *Review of Economic Studies* 56: 357:74.
- Carroll, Christopher D., 1994, "How does future income affect current consumption?", *Quarterly Journal of Economics*, February, 111-47.
- Carroll, Christopher D., and Miles Kimball, 1996, "On the concavity of the consumption function", *Econometrica*, LXIV, 981-92.
- Carroll, Christopher D., 1997a "Buffer-Stock saving and the Permanent Income Hypothesis", *Quarterly Journal of Economics*, 112 (1), 1-56.
- Carroll, Christopher D. 1997b. "Death to the log-linearized Euler equation! (And very poor health to the second order approximation." NBER Working paper No. 6298.
- Deaton, Angus, 1991, "Saving and liquidity constraints", *Econometrica*, 59(5), 1221-48.
- Deaton, Angus, 1992, *Understanding consumption*, New York, Oxford University Press.
- Dynan, Karen E., 1993, "How prudent are consumers?", *Journal of Political Economy*, 101(6), 1104-13.
- Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese, 1992, "Earnings uncertainty and precautionary saving", *Journal of Monetary Economics*, 30(4), 307-337.
- Hall, Robert, E., 1988, "Intertemporal substitution in consumption", *Journal of Political Economy*, 96(2),339-57.

- Kimball, Miles S., 1990, "Precautionary saving in the small and in the large", *Econometrica*, 58(1), 53-73.
- Kuehlwein, Michael, 1991, "A test for the presence of precautionary saving", *Economic Letters*, 37, 471-75.
- Laibson, David. 1997. "Hyperbolic Discount Functions and Time Preference Heterogeneity." Mimeo, Harvard University.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman. 1998. "Self Control and Saving for Retirement." *Brookings Papers on Economic Activity* 1: 91-196.
- Ludvigson, Sydney C. Forthcoming. "Consumption and Credit: A Model of Time-Varying Liquidity Constraints." *Review of Economics and Statistics*.
- Lusasrdi, Annamaria. 1993. "Precautionary Saving and Subjective Earnings Variance." VSB Progress Report No. 16. CentER, Tilburg U.
- Lusasrdi, Annamaria. 1996. "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets." *Journal of Business and Economic Statistics*. 14(1): 81-90.
- MaCurdy, Thomas E., 1982, "The use of time series processes to model the error structure of earnings in a longitudinal data analysis", *Journal of Econometrics*, 18, 83-114.
- Merrigan, Philip, and Michel Normandin, 1996, "Precautionary savings motives: An assessment from UK time series of cross sections", *The Economic Journal*, 106, 1193-08.
- Pischke, Jörn-Steffen, 1995, "Individual income, incomplete information, and aggregate consumption", *Econometrica*, 63(4), 805-40.
- Skinner, Jonathan, 1988, "Risky income, life cycle consumption, and precautionary savings." *Journal of Monetary Economics* 22, 237-255.
- Zeldes, Steve P., 1989, "Consumption and liquidity constraints: An empirical investigation", *Journal of Political Economy*, 97, 305-46.

Table 1 : Comparisons of linearized and true consumption growth functions

Time period	Baseline	$\sigma = .125$	$\phi=.30$	$\rho = 4$	$r = .05$	$\mu = .01$
Weighted averages of ratio of true slope to linearized slope (i.e. $(\rho^*+1)/(\rho+1)$ )						
18	0.865	0.943	0.799	0.771	0.840	0.869
21	0.860	0.944	0.796	0.775	0.840	0.870
24	0.858	0.944	0.799	0.777	0.837	0.866
27	0.857	0.945	0.796	0.770	0.837	0.866
30	0.858	0.947	0.793	0.770	0.836	0.867
33	0.857	0.947	0.796	0.766	0.835	0.865
36	0.852	0.950	0.796	0.765	0.831	0.861
39	0.854	0.945	0.788	0.750	0.830	0.858
42	0.846	0.940	0.788	0.745	0.830	0.853
45	0.840	0.930	0.774	0.735	0.824	0.847
48	0.830	0.915	0.757	0.731	0.817	0.838
51	0.821	0.899	0.737	0.704	0.809	0.828
54	0.808	0.882	0.643	0.673	0.799	0.814
57	0.763	0.839	0.680	0.643	0.762	0.787
Ratio of $\rho^*$ to $\rho$						
18	0.820	0.924	0.732	0.714	0.787	0.826
21	0.813	0.926	0.729	0.718	0.787	0.826
24	0.810	0.925	0.731	0.721	0.783	0.821
27	0.809	0.926	0.728	0.713	0.783	0.821
30	0.811	0.930	0.725	0.713	0.781	0.822
33	0.809	0.929	0.728	0.708	0.780	0.820
36	0.803	0.933	0.729	0.706	0.774	0.814
39	0.805	0.927	0.717	0.687	0.773	0.811
42	0.795	0.920	0.717	0.681	0.773	0.804
45	0.786	0.906	0.699	0.668	0.765	0.796
48	0.773	0.887	0.676	0.664	0.756	0.784
51	0.761	0.865	0.649	0.630	0.746	0.771
54	0.745	0.843	0.525	0.591	0.732	0.752
57	0.685	0.786	0.574	0.554	0.682	0.716

Notes: The parameters for the baseline case are:  $\sigma=.20$ ,  $\rho=3$ ,  $\phi=.44$ ,  $r=.03$  and  $\mu=.02$ . The method of computing weighted averages is described in the text. The results are from a model with  $T=60$ .

Table 2: Descriptive statistics and parametric assumptions  
(9,000,000 observations on 300,000 "individuals", 30 years per individual)

	<12 years school	12 years school	>12 years school
<u>Professional, technical and kindred</u>			
$\mu$	.	0.047	0.082
$\sigma$	.	0.153	0.147
$\phi$	.	0.159	0.288
% sample	.	1.40	14.46
Mean and (Std.dev.) of $\Delta\ln(c_{it})$	.	.0497 (.173)	.0796 (.225)
Mean and (Std.dev.) of $[\Delta\ln(c_{it})]^2$	.	.0325 (.051)	.0571 (.298)
<u>Managers (self-employed and employees)</u>			
$\mu$	.	0.027	0.056
$\sigma$	.	0.176	0.165
$\phi$	.	0.342	0.397
% sample	.	3.95	9.90
Mean and (Std.dev.) of $\Delta\ln(c_{it})$	.	.0294 (.137)	.0507 (.129)
Mean and (Std.dev.) of $[\Delta\ln(c_{it})]^2$	.	.0196 (.026)	.0338 (.112)
<u>Clerical and sales</u>			
$\mu$	0.038	0.038	0.053
$\sigma$	0.137	0.150	0.188
$\phi$	0.407	0.809	0.314
% sample	0.93	3.22	5.89
Mean and (Std.dev.) of $\Delta\ln(c_{it})$	.0250 (.129)	.0016 (.064)	.0581 (.187)
Mean and (Std.dev.) of $[\Delta\ln(c_{it})]^2$	.0172 (.057)	.0041 (.012)	.0383 (.071)
<u>Craftsmen and kindred</u>			
$\mu$	0.024	0.025	0.034
$\sigma$	0.196	0.170	0.171
$\phi$	0.329	0.341	0.273
% sample	6.29	12.92	6.48
Mean and (Std.dev.) of $\Delta\ln(c_{it})$	.0324 (.144)	.0264 (.132)	.0374 (.152)
Mean and (Std.dev.) of $[\Delta\ln(c_{it})]^2$	.0216 (.029)	.0180 (.024)	.0245 (.034)
<u>Operatives (transport and non-transport)</u>			
$\mu$	0.020	0.025	0.047
$\sigma$	0.171	0.170	0.177
$\phi$	0.432	0.366	0.299
% sample	6.93	9.83	3.25
Mean and (Std.dev.) of $\Delta\ln(c_{it})$	.0191 (.115)	.0254 (.129)	.0495 (.173)
Mean and (Std.dev.) of $[\Delta\ln(c_{it})]^2$	.0137 (.018)	.0173 (.032)	.0326 (.062)
<u>Laborers and service workers</u>			
$\mu$	0.046	0.048	0.050
$\sigma$	0.182	0.210	0.220
$\phi$	0.563	0.464	0.263
% sample	5.39	5.80	3.36
Mean and (Std.dev.) of $\Delta\ln(c_{it})$	.0342 (.147)	.0492 (.174)	.0631 (.197)
Mean and (Std.dev.) of $[\Delta\ln(c_{it})]^2$	.0228 (.065)	.0325 (.087)	.0247 (.062)

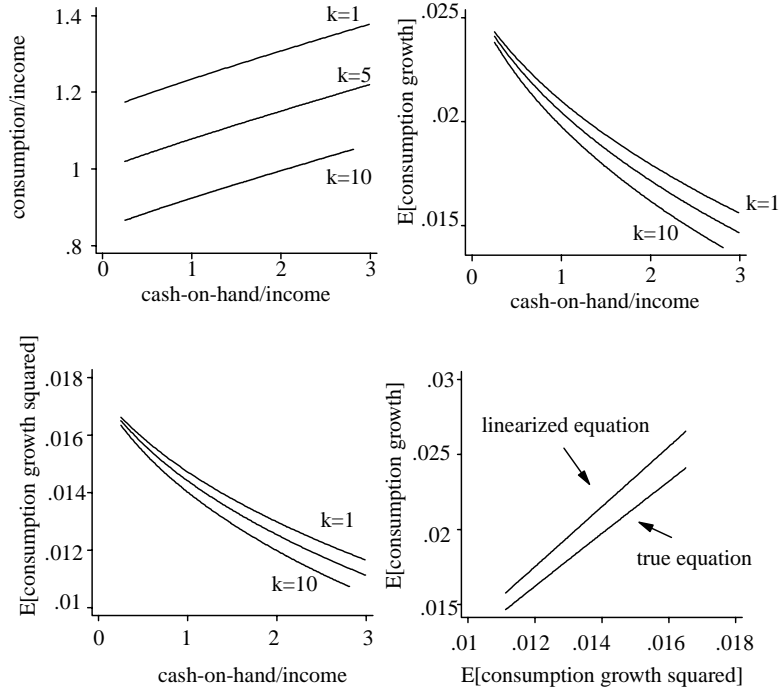
Table 3: Monte Carlo results

	Model A $cg_{it} = \beta_0 + \beta_1 cg_{it}^2 + \epsilon_{it}$			Model B $\frac{1}{T} \sum_{t=1}^T cg_{it} = \beta_0 + \beta_1 \frac{1}{T} \sum_{t=1}^T cg_{it}^2 + v_i$		
	full samp.	low wealth	high wealth	full samp.	low wealth	high wealth
Obs per regression # samples drawn	29,000 300	14,500 300	14,500 300	1,000 300	500 300	500 300
<u>OLS estimates</u>						
Average estimate of $\beta_0$	0.0145	0.0223	-0.0201	0.0225	0.0345	0.0038
Average estimate of $\beta_1$	0.968	0.864	2.461	0.688	0.514	1.299
Standard deviation of estimate of $\beta_1$	0.389	0.349	0.036	0.233	0.188	0.232
# rejections of $\beta_1=(\rho+1)/2$	300	300	300	300	300	299
<u>IV . Instrument Set 1: 15 occupation/education dummies</u>						
Average estimate of $\beta_0$	0.0016	0.0026	-0.0014	0.0016	0.0026	-0.0003
Average estimate of $\beta_1$	1.408	1.383	1.551	1.408	1.383	1.494
Standard deviation of estimate of $\beta_1$	0.128	0.147	0.174	0.128	0.151	0.189
# rejections of $\beta_1=(\rho+1)/2$	300	300	245	300	300	265
# times OID test is rejected	25	18	18	13	1	44
# rejections instruments are jointly insignificant	300	297	300	300	296	300
#rejections instruments are jointly insignificant in a regression of consumption growth cubed on instruments	250	183	300	265	117	216

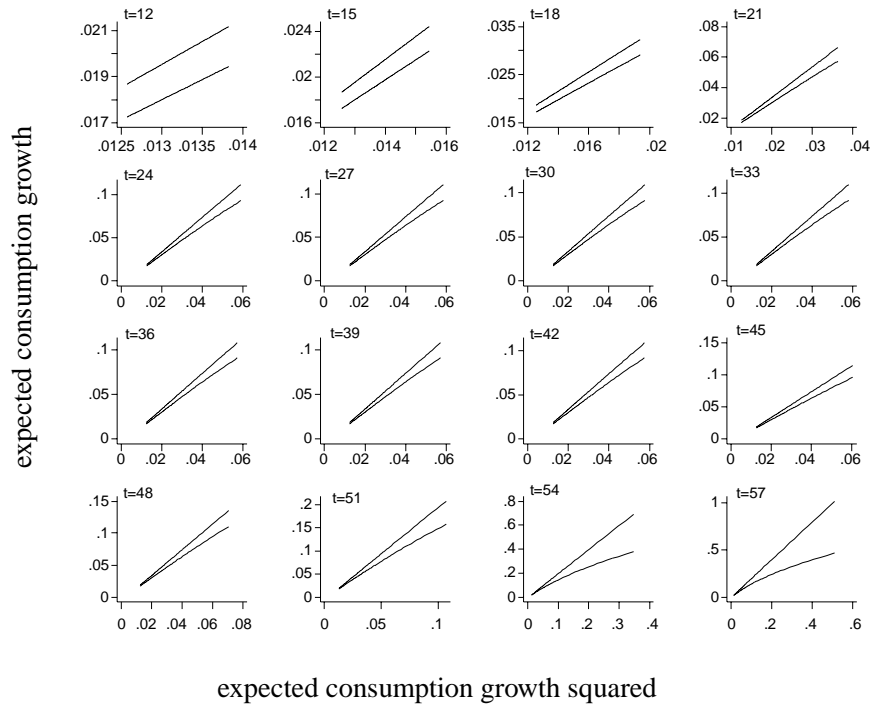
Table 3 (continued)

	Model A $cg_{it} = \beta_0 + \beta_1 cg_{it}^2 + \epsilon_{it}$			Model B $\frac{1}{T} \sum_{t=1}^T cg_{it} = \beta_0 + \beta_1 \frac{1}{T} \sum_{t=1}^T cg_{it}^2 + v_i$		
	full samp.	low wealth	high wealth	full samp.	low wealth	high wealth
<u>IV, Instrument Set 2: 6 occupation and 3 education dummies</u>						
Average estimate of $\beta_0$	0.0018	0.0031	-0.0016	0.0018	0.0029	-0.005
Average estimate of $\beta_1$	1.401	1.371	1.559	1.401	1.375	1.503
Standard deviation of estimate of $\beta_1$	0.130	0.151	0.189	0.130	0.153	0.190
# rejections of $\beta_1=(\rho+1)/2$ (5% level)	300	300	222	300	300	258
# times OID test is rejected (5% level)	27	11	24	27	4	34
# rejections instruments are jointly insignificant	300	298	300	300	298	300
#rejections instruments are jointly insignificant in a regression of consumption growth cubed on instruments	271	216	300	235	110	215
<u>IV, Instrument Set 3: 6 occ, 3 ed dummies and lagged assets/income</u>						
Average estimate of $\beta_0$	0.0046	0.0095	-0.0016	0.0017	0.0029	-0.0004
Average estimate of $\beta_1$	1.304	1.203	1.559	1.403	1.375	1.501
Standard deviation of estimate of $\beta_1$	0.159	0.181	0.189	0.130	0.155	0.190
# rejections of $\beta_1=(\rho+1)/2$ (5% level)	300	300	220	300	300	259
# times OID test is rejected (5% level)	144	101	20	21	5	28
# rejections instruments are jointly insignificant	300	300	300	300	298	300
#rejections instruments are jointly insignificant in regression of consumption growth cubed on instruments	299	300	300	255	113	209

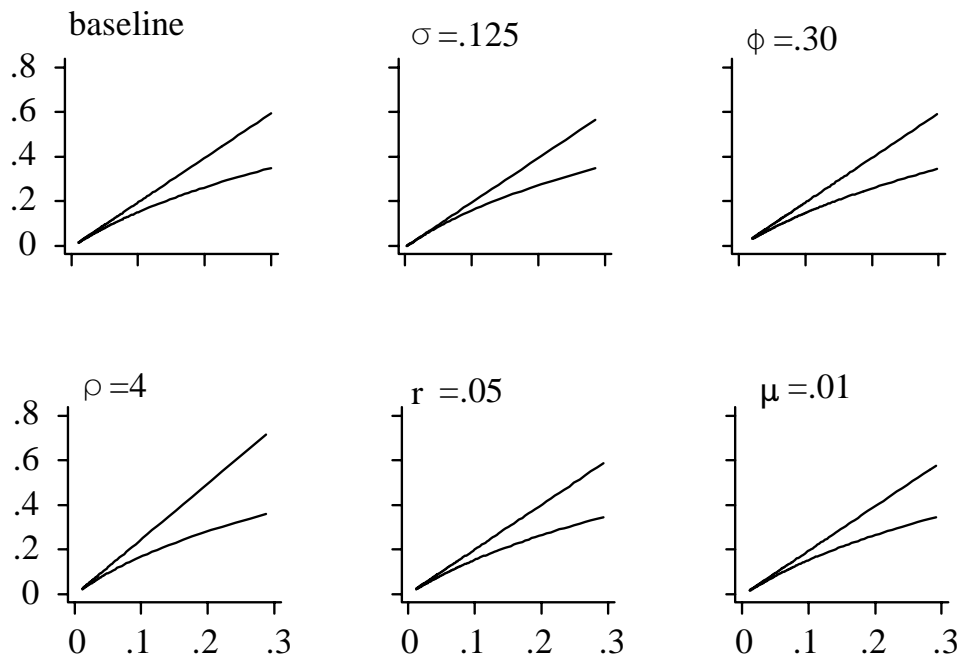





**Figure 1:** Baseline case,  $t=30$ , income states 1, 5 and 10




**Figure 2:** Baseline case,  $k=5$ , selected years




**Figure 3:** Deviations from baseline case,  $t=30$ , state=5




**Figure 4:** Distribution of OLS and IV estimates of  $\beta_1$ , Model A, full sample




**Figure 5:** Distribution of OLS and IV estimates of  $\beta_1$ , Model A, low-wealth sample



**Figure 6:** Distribution of OLS and IV estimates of  $\beta_1$ , Model A, high wealth sample




**Figure 7:** Distribution of OLS and IV estimates of  $\beta_1$ , Model B, full sample



**Figure 8:** Distribution of OLS and IV estimates of  $\beta_1$ , Model B, low wealth sample





**Figure 9:** Distribution of OLS and IV estimates of  $\beta_1$ , Model B, high wealth sample