APPROXIMATION BIAS IN LINEARIZED EULER EQUATIONS

Sydney Ludvigson and Christina H. Paxson

Federal Reserve Bank of New York
Research Paper No. 9712

March 1997

Approximation Bias in Linearized Euler Equations

Sydney Ludvigson*
Christina H. Paxson**

March 5, 1997

Abstract

This paper concerns pitfalls associated with the use of approximations to dynamic Euler equations. Two applications of the approximations are notable. First, tests for precautionary saving motives typically involve regressing consumption growth on uncertainty in expected consumption growth. The parameter estimates are used to measure the strength of precautionary motives, which is also related to the coefficient of relative risk aversion. Another application estimates the sensitivity of consumption growth to the expected real interest rate, with the coefficient on the latter equal to the intertemporal elasticity of substitution in consumption, often the inverse of the coefficient of relative risk aversion. The two literatures yield very different estimates of how prudent or risk averse consumers are or, alternatively, how willing they are to substitute consumption over time. We investigate one possible reason for these apparently contradictory results: both methods of estimation rely on linear *approximations* of Euler equations. We demonstrate that biases associated with these approximations can be substantial, and that the direction of the biases is consistent with the divergent estimates found in the literature.

* Federal Reserve Bank of New York, Domestic Research Function - 3F, 33 Liberty Street, New York, NY 10045, sydney.ludvigson@frbny.sprint.com.
** Department of Economics and Woodrow Wilson School, 219 Bendheim Hall, Princeton University, Princeton, NJ 08544, cpaxson@wws.princeton.edu.
We thank Bo Honore, Angus Deaton, Steve Zeldes, and workshop participants at the NY Federal Reserve and Princeton University for helpful comments. Paxson received support from the National Science Foundation for this work. The opinions expressed in this paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors and omissions are the responsibility of the authors.

I. Introduction

This paper concerns pitfalls associated with the use of approximations to dynamic Euler equations. There are a variety of economic applications that rely on approximations to dynamic Euler equations when no closed form solution for the optimal consumption rule exists. In this paper we focus on two of these applications. First, there is a large and growing empirical literature on precautionary saving that examines how consumption growth and saving behavior are affected by uncertainty. Tests for precautionary saving motives typically involve regressing consumption growth on measures of uncertainty in expected consumption growth, the idea being that future uncertainty will depress current consumption and raise consumption growth. This 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 precautionary saving motives, where in the absence of precautionary motives future uncertainty should not affect consumption growth.

The second application concerns the effects of movements in the expected real interest rate on consumption growth. Here, the major assumption is that the interest rate is stochastic. It can be shown by taking a second-order Taylor expansion of the Euler equation that consumption growth is approximately linearly related to the conditional expected real interest rate, with the coefficient on the interest rate equal to the intertemporal elasticity of substitution or the inverse of the coefficient of relative risk aversion for popular forms of the utility function[1]. In this case,

---

[1]Recently, the asset pricing literature has developed more general forms of the utility function which break the link between the coefficient of relative risk aversion and the intertemporal elasticity of substitution. See, for example, Epstein and Zin (1991) and Weil (1989).

regressions of consumption growth on expectations about future interest rates are used to derive measures of the intertemporal elasticity of substitution and the degree of relative risk aversion.

Most of the empirical work in both of these literatures is based on linear approximations of Euler equations. However, the two literatures yield very different results. The estimated effects of consumption uncertainty on consumption growth are typically small, indicating that precautionary motives are weak or nonexistent. For utility functions characterized by decreasing absolute risk aversion, these results also imply implausibly low levels of relative risk aversion. However, the estimated effect of the expected real interest rate on consumption growth is also typically small, indicating a small intertemporal elasticity of substitution or a very *high* degree of relative risk aversion. More specifically, given that the within-period utility function is isoelastic, such that $u(C_t) = \frac{C_t^{1-\rho}}{1-\rho}$, the first literature yields estimates of ρ that are generally below 1.3, and are often insignificantly different from zero. The second literature typically yields estimates of ρ that range from 3 to 10, and are sometimes even higher.

In this paper we investigate one possible reason for these apparently contradictory results Specifically, both methods of estimating ρ rely on linear *approximations* of Euler equations. However, if the Euler equations are very nonlinear, then these approximations may be quite poor, and this can result in biased parameter estimates. We demonstrate that these biases can be substantial, and that the direction of the biases is consistent with the divergent estimates of ρ found in the literature. Our approach is to compute, numerically, true consumption functions for each model. We can then contrast the actual relationships between consumption growth, uncertainty in consumption growth, and the expected real interest rate, with the relationships implied by the linear approximations.

Section II discusses in more detail how approximations to Euler equations have been used in previous literature, and why estimates based on these approximations can be biased. Section III describes our methods for computing consumption functions, and shows the results of these computations. Section IV concludes.

II. Approximations to Euler Equations

We start with a simple 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 δ. Sub-utility functions in each period are identical and isoelastic-elastic. The maximization problem is summarized as:

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

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

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

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

We consider two models, each of which involves imposing restrictions on R_{t+1} and y_{t+1} in (2). First, following most of the literature on precautionary saving, we assume that the real interest rate in not stochastic and is fixed at $r = R - 1$. The only uncertainty consumers face is in labor income, which fluctuates from period to period. Second, following most of the macro literature on consumption and real interest rates, we assume that the real rate of return is stochastic but that labor income is fixed at y. There are two strains within the macro literature. One strain focuses on portfolio choice, so that A_t may be thought of as total wealth derived from a "market" portfolio of assets and including the (risky) returns to human capital accumulated by the individual. In this case "labor income" includes only (non-risky) income, y, from non-accumulated assets. For example, Merton (1969), Samuelson (1969), and Campbell (1993) among others, use extreme versions of this model, in which y is set to zero and all income is modeled as asset income. The second strain focuses on estimating the intertemporal elasticity of substitution in consumption and assumes that consumers receive labor income in every period, but can trade off consumption today for consumption tomorrow by accumulating and decumulating a single, liquid asset, A_t. Since this literature focuses on estimating a linearized version of the Euler equation which does not explicitly depend on labor income, it is not necessary to make any assumptions about the stochastic properties of earnings. Examples in this vein are Hall (1988), Campbell and Mankiw (1989), Deaton (1992), and Atanasio and Weber (1993). In our model with risky interest rates, we follow the second strain of macro literature and assume a single asset with risky return gross return, R_t, but fixed labor income, y.

A. Risky labor income

For the case of a certain interest rate and risky labor income, the Euler equation associated with utility maximization is:

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

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

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

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, what is commonly done is to 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:

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

where $1+\rho$ equals the coefficient of relative prudence $-\frac{u'''(C_t)C_t}{u''(C_t)}$, defined by Kimball (1990), which equals zero if there are no precautionary saving motives. The error term, ν_t, is composed of an additive series of moments of consumption growth. Specifically:

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

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 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. 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.

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 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 ρ (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 ρ 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 ρ, in the range of .78 to 1.33 for Merrigan and Normandin. Kuehlwein estimates a higher value for ρ, equal to about 0.25, but it is not statistically different from zero.

All of this work relies on IV estimation, and the identifying assumption is that the instruments—typically variables which measure industry, occupation, education, and employment status—are *not* correlated with the error term in (5). However, as shown in (6), the error v_t consists of terms that contain third and higher moments of consumption growth. It seems implausible that if the instruments are correlated with uncertainty in consumption growth, they would be uncorrelated with skewness, kurtosis, and higher moments of consumption growth.[2] Furthermore, it is likely that the approximation error v_t is correlated with uncertainty in consumption growth, and this will produce biased parameter estimates if (5) is estimated using ordinary least squares. 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 yield estimates of ρ 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 individual's true objective function, the bias would still exist—but is instead solely the result of approximation error.

B. Risky interest rates

When there is stochastic variation in the real interest rate, the Euler equation associated with utility maximization takes into account expectations about future interest rates:

---

[2]Most studies (e.g. Dynan, 1993; Merrigan and Normandin, 1996) report the results of overidentification (OID) tests, and in most cases these tests pass. However, it is not clear whether or not the tests pass because many of the instruments have little explanatory power in the first stage regression. Furthermore, it should be noted that OID tests are not equivalent to tests for the validity of identifying restrictions; the tests can, in principle, pass even though the instruments are correlated with some unobserved component of the error term.

$$u'(C_t) = (\frac{1}{1+\delta})E_t[(1+r_{t+1})u'(C_{t+1})], \quad (7)$$

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

$$C_t^{-\rho} = (\frac{1}{1+\delta})E_t[(1+r_{t+1})C_{t+1}^{-\rho}]. \quad (8)$$

As before, equation (8) can be transformed into a consumption growth equation by taking a second-order Taylor expansion, in this case of next period's marginal utility times the gross return on assets around the points $C_{t+1} = C_t$ and $r_{t+1} = 0$. The resulting consumption growth equation is:

$$E_t\left[\frac{C_{t+1}-C_t}{C_t}\right] = \frac{1}{\rho}\left[\frac{\delta}{1+\delta}\right] + \frac{1}{\rho}E_t r_{t+1} + (\frac{1+\rho}{2})E_t\left[\left(\frac{C_{t+1}-C_t}{C_t}\right)^2\right] - cov\left(\frac{C_{t+1}-C_t}{C_t}, r_{t+1}\right) + v_t, \quad (9)$$

where the error term v_t contains higher-order moments and co-moments of consumption growth and next period's real interest rate. Note that the coefficient on the expected real interest rate is defined as the intertemporal elasticity of substitution in consumption ($IES \equiv 1/\rho$). An alternate method of deriving a estimable consumption-growth equation is to assume that consumption and asset returns are jointly lognormal, see for example Campbell and Mankiw (1989) and the discussion in Chapter 2 of Deaton (1992). The resulting equation is a logarithmic version of (9):

$$E_t[\ln(C_{t+1}) - \ln(C_t)] = -\frac{\delta}{\rho} + \frac{1}{\rho}E_t r_{t+1} + \tfrac{1}{2}\rho\, var_t[\ln(C_{t+1}) - \ln(C_t) - \frac{r_{t+1}}{\rho}], \quad (10)$$

which also implies that consumption growth (now measured as the change in logs) is linearly related to the expected real interest rate, again with coefficient equal to $1/\rho$.

Equation (9) and its variant (10) have typically been estimated using time-series data. The growth rate of aggregate consumption is regressed on next period's interest rate, which is instrumented with variables known at time t to account for the fact that it is the conditional *expected* future interest rate that affects consumption growth. The right-hand-side variables other then the expected real interest rate are included in the error term. The resulting estimated coefficient implies a value for $1/\rho$, or the IES.

Though equations (9) and (10) relate consumption growth to the expected return on a single, risky asset, in practice there are many different types of assets and a question arises as to which asset's returns to use in estimation.[3] The literature has made use of a variety of returns to estimate the sensitivity of consumption growth to the expected real interest rate (the IES) in equations (9) or (10). Overall, the estimates support the idea that ρ is much larger than the literature on precautionary saving would suggest. For example, Campbell and Mankiw (1989) use the real three month t-bill rate and estimate a value of $1/\rho$ equal to 0.276, implying ρ equals 3.6. Hall (1988) uses several rates of return, the real three month t-bill rate, the real return on an indexed stock portfolio, and the regulated pass book interest rate on savings accounts. His results indicate that the IES is between .03 and 0.1, implying that ρ ranges from 10 to the implausibly high value of 33. Attanasio and Weber (1993) use the deflated interest rate on building society deposits (Britain) and estimate a statistically significant IES parameter to be about 0.35, implying

[3]Note that in a model of optimal portfolio choice with isoelastic utility and several different assets, (9) and (10) will hold for each asset in the portfolio.

ρ is approximately 3. Wirjanto (1994) uses the three month prime corporate paper rate (Canada) and estimates a statistically significant, but nevertheless relatively low value for the IES, equal to about 0.25, or ρ equal to 4. Campbell, Lo, and MacKinlay (1997, chapter 8) use a real indexed stock return and the real six month commercial paper rate, finding in both cases that the estimated value for the IES is insignificantly different from zero and that point estimates are often negative. Deaton (1992) provides a useful summary of results.

Two points deserve mention. First, as for the model of precautionary saving discussed above, approximation error may result in biased estimates of the intertemporal substitution elasticity. The second and higher-order moments and co-moments of consumption growth and the expected real interest rate need not, in general, be orthogonal to the expected real interest rate, and the variables used as instruments for the expected real interest rate (typically lagged interest rates) may also forecast these higher moments. Our work in the next section examines the direction and size of these biases.

Second, existing research makes clear that there are other reason why estimates of (9) and/or (10) may yield biased estimates of ρ. One issue that has been discussed is aggregation bias: even if (9) or (10) is a good description of individual behavior, it may be unreasonable to think that these equations should describe the relationship between growth in *aggregate* consumption and the real interest rate. Attanasio and Weber (1993) provide evidence that when disaggregated data are used to estimate (10), the estimate of ρ is a much more reasonable 1.3. It may be that aggregation is a much more important source of bias than is the linear approximation of Euler equations.

III. Numerical Solutions to Consumption Functions

To assess the extent to which approximation bias 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 ρ and the stochastic process that governs the evolution of income and the interest rate.

A. Risky labor income

We start with the first model discussed above, in which the interest rate is fixed but there is uncertainty in labor income [4]. The first step is to choose a utility function and a stochastic process for income. We use an isoelastic utility function with a variety of values for ρ. We assume that income growth follows a first-order moving-average process with the general form:

$$\ln(y_t) = \ln(y_{t-1}) + \mu + \epsilon_t - \phi\epsilon_{t-1}. \tag{11}$$

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 (1996). Though 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.

---

[4] 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.

In the simpler case of independently distributed income, the second step would be to solve for consumption in each period as a function of cash-on-hand, i.e. income plus assets (positive or negative) carried in from last period. Serial correlation in income complicates matters, since it means that there will be two state variables (cash-on-hand and lagged income) rather than one. 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 both, 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.

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, ϵ_t. Let θ_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:

$$\theta_t(w_t,\epsilon_t)^{-\rho} - \beta \int \theta_{t+1}(\,[1+r][w_t-\theta_t(w_t,\epsilon_t)]z_t^{-1} + 1\,,\epsilon_{t+1})^{-\rho} z_{t+1}^{-\rho}\, dF(\epsilon_{t+1}) = 0\,, \qquad \textbf{(12)}$$

where β equals (1+r)/(1+δ).

The Euler equations are solved via backwards recursion, starting with a terminal time period T. We choose a grid of 200 values of w_T, and solve for a set of corresponding values of θ_T. Although it will generally be the case that θ is a function of the income innovation ϵ 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 equals cash-on-hand, so that θ_T equals w_T for all values of ϵ_T. Given this solution for θ_T, one can solve for the values of θ_{T-1} that satisfy the Euler equation for a grid of possible values of w_{T-1} and ϵ_{T-1}. Solutions for earlier periods can be found by working backwards, i.e. solving for θ_t as a function of w_t and ϵ_t, given the solutions for θ_{t+1}.

Two technical points deserve mention. First, the distribution of the error ϵ_t is specified as a discrete 10-point approximation to a normal distribution, which considerably speeds computation time. Second, and more importantly, the grid of values for the ratio of cash-on-hand to income, w_t, must be chosen with considerable care, and must be allowed to differ from period to period. The problem is that the solutions to θ_t in any period imply a possible range of values of w_{t+1}: one cannot solve for values of θ_t which, if chosen, can result in values of w_{t+1} that were not included in the grid range when solving for θ_{t+1}. To see this problem more clearly, consider the identity that relates cash-on-hand in t+1 to cash-on-hand and consumption in t (where all variables are expressed as ratios to income):

$$w_{t+1} = (1+r)[w_t - \theta_t]z_{t+1}^{-1} + 1 . \qquad \textbf{(13)}$$

Since we work backwards in time, when we solve for θ_t as a function of w_t, the range of grid values for w_{t+1} has already been set in the previous iteration. Let the maximum and minimum values of w_{t+1} be $w_{t+1}^{\max}$ and $w_{t+1}^{\min}$. Note that because there are 10 possible values of the shocks ϵ_t, there are 100 possible values of z_{t+1}, the ratio of income in t+1 to income in t. We must choose a range for w_t such that for all 100 combinations of ϵ_{t+1} and ϵ_t:

$$w_{t+1}^{\min} \le (1+r)[w_t - \theta_t(w_t,\epsilon_t)]z_{t+1}^{-1} + 1 \le w_{t+1}^{\max} \tag{14}$$

It is impossible to solve analytically for the maximum feasible range of w_t. Instead, we must seek a range for w_t by trial and error. Note from (14) that the wider the range of w_t, the more likely w_{t+1} in (13) will fall outside the range specified in (14). Our approach is to choose a (generous) initial range of w_t, see if it is possible to solve the Euler equation without violating (14), and then contract the range of w_t if necessary. A problem with this general procedure is that the range for the cash-on-hand to income ratio generally shrinks as one goes back in time, and this limits the number of "years" for which the Euler equations can be solved. In what follows we present results for 6 years (i.e. up to T-6). Although this may seem short, we show that the consumption functions for which we solve do not vary much from period to period prior to year T-2.

An alternative procedure is to use the same range of values for w_t in each period. However, in this case one must devise some way to "fill in" values of θ_{t+1} and w_{t+1} which are required to solve the Euler equation, but which have not been solved for; i.e., the condition in (14) fails. A common procedure is to use linear extrapolation to impute the values of $\theta_{t+1}(w_{t+1},\epsilon_t)$ and w_{t+1} that lie outside of the grid range $w_{t+1}^{\min}$ to $w_{t+1}^{\max}$. The advantage of using linear extrapolation is that the grid range need not shrink as one goes back in time, and it is possible to compute consumption functions for many more time periods. The disadvantage is that linear extrapolation introduces errors into the solutions, and these errors are compounded as one works backward in time. We experimented with computing consumption functions over a fixed grid range using linear extrapolation, and found the errors induced by extrapolation significantly altered the results. We concluded that the benefit of having consumption functions for more "years" was not worth the costs due to extrapolation error. We discuss this issue further below.

Figure 1 shows results from what we call our "baseline case." In this case, we have set the parameters of the income equation (11) to μ equal to .02, ϕ equal to .444, and the standard deviation of ϵ (denoted σ) to .25. The parameters values for ϕ and σ are from MaCurdy (1982) and Pischke (1996). The estimated values for σ and ϕ may be biased by measurement error in recorded income; therefore we vary the values of these parameters to reflect this possibility in several other cases discussed below. The baseline interest rate r is set to .03, the rate of time preference, δ, is set to .05, and the parameter ρ of the isoelastic utility function is set to 3. We used these parameters to solve for six years of consumption functions, excluding the last period of life, T, in which consumption equals cash-on-hand.

Figure 1 shows characteristics of the consumption function for period T-6, the "earliest" year for which we have a solution. The top left-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:

$$\frac{E[\theta_{t+1} z_{t+1} | w_t, \epsilon_t] - \theta_t(w_t, \epsilon_t)}{\theta_t(w_t, \epsilon_t)}, \tag{15}$$

which is easily computed for each of the possible values of w_t and ϵ_t. Note that as many as ten functions could be graphed, one for each of the ten possible values of ϵ_t. To avoid clutter, we graph only three. The line marked "k=1" denotes the graph for the lowest value of ϵ, "k=5" is the fifth-lowest value, and "k=10" is the highest. In the same way, we can easily compute expected squared consumption growth, which is shown plotted against cash-on-hand to income in the upper right-hand panel of Figure 1.

Two features of the graphs in the top panels of Figure 1 illustrate the main properties of the consumption functions. First, individuals with low cash-on-hand (relative to income) have both higher expected consumption growth, and higher expected squared consumption growth. These results are consistent with the findings of Carroll (forthcoming) and Carroll and Kimball (1996): because poor consumers have a lesser ability to buffer shocks to income, the conditional variance of consumption growth rises and precautionary motives work to depress consumption and increase its rate of growth. Second, for any given cash-on-hand to income ratio, consumption growth is higher for lower values of ϵ_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 decumulating 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 displayed in the top two panels of Figure 1, we are now in a position to compare their relationship with that implied by the Taylor expansion (5). These comparisons are made in the bottom two panels of Figure 1. The bottom left-hand panel graphs the relationship between expected consumption growth and expected squared consumption growth. There are actually ten different functions relating these two variables, one for each value of ϵ, 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. The consumption growth equation implied by the Taylor expansion (5), indicates the function graphed "should" have an intercept equal to $\frac{1}{\rho}[\frac{r-\delta}{1+r}]$ and a slope of $(1+\rho)/2=2$ given the baseline parameters. This linearized equation is shown in the lower left hand panel. The lower right-hand panel shows how much the linear approximation overstates the actual relationship between expected consumption growth and expected consumption growth squared, by plotting the ratio of the true slope implied by the optimal consumption policy, to $(\rho+1)/2$, the slope implied by Taylor expansion. To simplify notation in what follows, let the true slope be denoted as $(\rho^*+1)/2$, although it should be kept in mind that ρ^* is not a fixed number, and varies with the cash-on-hand to income ratio.

The bottom two panels of Figure 1 give a visual impression of the bias induced by linearization. The left panel shows that the slope of the true function relating expected consumption growth to the expected squared consumption growth is lower than that implied by the equation based on linear approximation. The right panel indicates that the approximation bias in the slope can be substantial: for a reasonable range of cash-on-hand to income, the true slope is found to be between 70% and 80% of what would be inferred from equation (5). The bias in ρ is even larger in percentage terms. For example, if the ratio of $(\rho^*+1)/2$ to $(\rho+1)/2$ is .75 and ρ equals 3, then ρ^* is equal to 2, only 67% of ρ.

An important aspect of these results is that the bias also varies with wealth, and is more severe for less wealthy households. It is therefore not surprising to find that estimates of ρ 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 ρ downward for less wealthy consumers, and, consistent with this hypothesis, they report lower values for less wealthy households. Note, however, that these findings are also consistent with the direction of bias shown in Figure 1, 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. We now ask how such bias changes over time.

Figure 2 graphs the ratio of the true slope (in the relationship between expected consumption growth and its variance) to that implied by the Taylor expansion for k=5 for all six periods, again using the baseline parameters. As the figure shows, the bias decreases as we begin to move backward from the final period, T. For example, at T-1 the true slope ranges from 40 to 60 percent of what is implied by the Taylor approximation, depending on the level of cash-on-hand. At T-2 this ratio is roughly between 60 and 70 percent, and by the time we reach T-3 the bias in the slope levels off to between 70 and 80 percent where it remains as we move further back in time. In other words, the optimal consumption functions "converge" quite rapidly.

It should be noted that, although we graph the results for each of the years only for values of the cash-on-hand to income ratio of .5 to 1.5, we actually computed consumption functions for a much wider ranges of this ratio for the periods closer to T. For example, for year T-1 the functions were computed for cash-on-hand to income ratios that ranged from .0005 to 8.8. To avoid linear extrapolation, this range was narrowed as we worked backward, so that the ratio ranged from .528 to 1.6 by T-6. If we had wanted to solve for many more "years" (without linear extrapolation), we would have had to begin, in T, with a much wider range of cash-on-hand to income ratios.

We examined how different the results of Figures 1 and 2 would have been had we used linear extrapolation to solve for the consumption functions. Figures 3 and 4 correspond to Figures 1 and 2, only they were derived by restricting the range of cash-on-hand to income to .5 to 1.5 in *all* of the time periods, and using linear extrapolation when necessary. Figure 3 demonstrates that linear extrapolation has pronounced effects on the results. Note that the consumption functions themselves do not seem to be greatly affected by extrapolation error. This can be seen from the top two panels which do not appear too different from those in Figure 1. However, linear extrapolation clearly affects the *relationship* between expected consumption growth and expected squared consumption growth, shown in the bottom two panels. For example, the results with extrapolation imply that for high values of cash-on-hand relative to income, and for "good" income draws (k=10), ρ is actually biased up rather than down. But, when the cash-on-hand to income ratio and the income innovation are high, the consumer will choose to enter the next period with a high level of assets, and is therefore more likely to have a cash-on-hand to income ratio that is higher than the maximum value of cash-on-hand to income that has been set. Thus, these "results" reflect the effects of extrapolation rather than a genuine feature of the model. Likewise, Figure 3 indicates that for low values of cash-on-hand to income and "bad" income draws, ρ is biased down by even more than is shown in Figure 1. Again, however, this is solely the result of errors introduced by extrapolation. Figure 4 (which, like Figure 2, shows results for k=5) indicates that these errors become "compounded" as one works backwards through time periods. For T-1, the results from Figure 2 (without extrapolation) and Figure 4 (with extrapolation) are nearly identical. By T-6 there are pronounced effects of extrapolation at the "edges" of the functions. For all of the results that follow, we do not rely on linear extrapolation.

How sensitive are our results to changes in the baseline set of parameters? Figures 5 through 11 show results that correspond to those in Figure 2, but with changes from the baseline case. Again, we graph results for k=5. It should be kept in mind, however, that within each year the extent of bias depends not only on the cash-on-hand to income ratio w, but also on the innovation to income ϵ (with k=1 to 10 representing the ten possible values of ϵ). To summarize the results shown in the figures, we compute for each time period the weighted average of the true slope $(\rho^{*}+1)/2$ to the linearized slope $(\rho+1)/2$. The weights are constructed in the following way. First, we simulate the consumption and saving decisions of 5000 consumers: each consumer is assumed to begin in T-6 with cash-on-hand equal to income, so that w_{T-6} equals 1. We then draw values of ϵ_{T-5} through ϵ_{T-1} for each consumer, and use the policy functions for the appropriate year to solve for θ and for the cash-on-hand to income ratio w in each year. Second, we use the simulated data to estimate the joint density of (w, ϵ) for each year from T-5 to T-1.[5] This density function is used to compute the weighted average of the ratio of the true to linearized slope. The results are shown in Table 1.

The results of Table 1 and the corresponding figures are summarized as follows. First, the extent of bias is very sensitive to σ, the standard deviation of ϵ. When σ is reduced from .25 to .125, the ratio of the true slope to that implied by equation (5) rises, from an average of .75 to .97 in T-5, implying an increase in ρ^{*}/ρ from .67 to .96 (see columns 1 and 2 of Table 1, and Figure

[5] The density is computed nonparametrically, using a quartic kernel with a bandwidth of 1.

5). An implication is that estimates of ρ 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, the bias is also very sensitive to ϕ, the moving average parameter. We lowered ϕ from .44 to .20. The results shown in Table 1 and Figure 6 indicate that this increases the extent of downward bias, so that the ratio of the slopes falls from .76 in the baseline to .68. Decreases in ϕ 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 ϕ have effects similar to increases in σ.

The results discussed above indicate that a smaller value for σ will lead to less downward bias in ρ, whereas a smaller value for ϕ will lead to more bias. It should be noted that empirical estimates of both parameters (σ and ϕ) will tend to be biased upward if there is measurement error in recorded income. For example, suppose that σ and ϕ are estimated at 0.25 and 0.44, respectively, from (11). If, due to the presence of measurement error, the true standard deviation of the income innovation is only 75% of what is estimated (so that σ is really 0.19 instead of 0.25), the true value of ϕ would be 0.26 instead of 0.44.[6] In Figure 7, we show the case with

---

[6]To see this, consider the following equation for the log difference in income: $\Delta \ln y_t = \mu + \eta_t - \lambda\eta_{t-1} + v_t - v_{t-1}$, where the last two terms are the first difference in i.i.d. measurement error and η is the true innovation to the growth in income. If income growth is estimated as in (11) with $\Delta \ln y_t = \mu + \epsilon_t - \phi\epsilon_{t-1}$ yielding estimates of the variance of ϵ and ϕ equal to 0.25 and 0.44 respectively, then by equating variances and covariances across the two equations, the variance of η, σ^2, is inversely related to the true moving average parameter, λ: $\sigma^2 = (0.01932)/(1-\lambda)^2$.

ϕ=0.26 and σ=0.19. The results indicate that there is somewhat less downward bias than in the baseline case, but it is clear that the decline in ϕ offsets some of the effects of the decline in σ.

Figures 8 and 9 show the effects of changing the parameter ρ. When ρ is reduced from 3 to 2 (Figure 8), the amount of bias falls, so that the ratio of the true slope to the linearized slope averages .87 in T-5. A value of ρ of 4 (Figure 9) reduces the average ratio to .66. Note this implies the value of ρ which would be inferred from the linearized regression, ρ^*, is only 57 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. Figures 10 and 11 show the effects of changing the parameters r and μ. Our 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 μ has little effect.

B. Risky Interest Rates

In this section we analyze the relationship between expected consumption growth and the expected real interest rate, when the latter is stochastic but labor income is deterministic.[7] We assume that the interest rate follows a first-order autoregressive process, which seems to fit the data reasonably well for a variety of asset returns (Campbell, Lo and MacKinlay, 1997; Fama and French, 1988). The gross return $R_{t+1} \equiv 1+r_{t+1}$ is assumed to follow:

[7]As in the model with risky labor income, the model with risky interest rates does not in general have a closed form solution. However, for explicit assumptions about asset returns and labor income an analytical solution can be derived. For example, Merton (1969) and Samuelson (1969) both show that if R is identically and independently distributed, and there is no labor income risk, the optimal consumption will be proportional to wealth in every period. We do not focus our analysis on this special case because we want to analyze the relationship between consumption growth and time-varying, conditional expected returns. Nevertheless, as a robustness check on our numerical solutions, we verified that the computed solution was the same as the analytical solution derived in this special case.

$$R_{t+1} - g = \alpha(R_t - g) + e_{t+1}, \tag{16}$$

where g is the mean return and $0 < |\alpha| < 1$.We will assume the error term is normally distributed, with mean zero and constant variance, σ_e and as before, we make a discrete approximation to the underlying distribution for e. We solve the model for a variety of values of g, α and σ_e. These values are taken from estimates of (16) for several different types of assets (commercial paper, stocks, and t-bills), the returns of which have been used in past research when estimating the effect of interest rates on consumption growth as in (10).

When solving for the optimal consumption functions, there are two methodological differences from the case of risky labor income discussed above. First, though the consumption functions will continue to depend on cash-on-hand, the latter now evolves in a stationary way since the driving process in (16) is no longer nonstationary. As a result, we can solve directly for the level of consumption as a function of the level of cash-on-hand, rather than having to define the Euler equation in terms of ratios of variables to current income. Second, the interest factor R_t is now a state variable, since autocorrelation in asset returns implies that R_t forecasts R_{t+1}.

With these changes, the solution proceeds as before, by backward recursion on the following set of dynamic equations:

$$C_t(x_t,R_t)^{-\rho} - \beta\int C_{t+1}(R_{t+1}[x_t - C_t(x_t,R_t)] + y, R_{t+1})^{-\rho} dF(R_{t+1}|R_t) = 0, \tag{17}$$

where y is deterministic labor income. The term $dF(R_{t+1}|\, R_t)$ is the density of next period's gross return on assets conditional on this period's return. We specify the joint density of R_{t+1} and R_t as a 25 by 25 discrete approximation to a bivariate normal, so that conditional on any value of R_t

there are 25 possible outcomes for R_{t+1}. The choice of 25 points represents a compromise. Using a large number of values for the gross return is more realistic, and allows one to examine how the properties of the consumption function differ across a larger set of returns. However, one must solve for separate consumption functions for each possible value of the state value R_t, so that the more points one uses the longer the computations take.

To set values for the model's underlying parameters, we choose a range of values for α, σ_e and ρ and leave the other parameters unchanged from our baseline case in the last section (results do not change qualitatively when varying these other parameters). Using time series data on the six month commercial paper rate and the three month t-bill rate, and using the implicit price deflator to estimate expected inflation, we estimated α at about .9 for both assets, with σ_e differing across the two returns and equal to 0.01, and 0.05, respectively. Alternatively, if R is the return on an indexed stock portfolio such as the Standard and Poor 500, Fama and French (1988) report that returns are slightly negatively autocorrelated at annual frequencies, and so we also consider the case when $\alpha = -.1$, with σ_e much is higher at 0.15 (see Campbell, Lo, and MacKinlay, 1997). Finally, we allow ρ to take on the values of 3, 6, and 8. For this model we define our baseline case as: ρ=3, α=.9, σ_e=0.05, and g=1.04. Deterministic labor income is normalized to 1.

As before, our strategy is to compute numerical solutions to the dynamic programming problem in (17), which yields the conditional expected value of consumption growth over a grid of values for x_t and R_t. Since C_t is a function of x_t and R_t, we can compute the true relationship between consumption growth and the expected real interest rate from 25 values of C and R, for any given level of cash-on-hand, x. The ratio of the this true slope (which we define to be $1/\rho^*$)to what would be implied from the Euler equation in (9) or (10) (equal to $1/\rho$) is given in Figure 12, for baseline parameter values, for each of the time periods. It should be kept in mind that we could have actually graphed 200 lines in each graph, one for each of the 200 values of x_t. To avoid clutter, we average the functions over all values of x_t from .5 to 5. Thus, the figure masks the fact that there is actually a non-monotonic relationship between the degree of upward bias in ρ and the level of cash-on-hand: as wealth increases, the bias rises at first, but then drops. However, for any given value of $E_t\,[r_{t+1}]$, the ratio ρ^*/ρ only varies by about 5 percent across values of cash-on-hand relative to income ranging from .5 to 5.

Figure 12 displays several notable qualities. First, for most reasonable values of the expected real interest rate, the linearized equation will overstates the parameter ρ, or *under*states the IES (equal to $1/\rho$.) In general, it is only for a few large negative values of the real interest rate that the bias works to overestimate the slope. Moreover, the upward bias in ρ increases as the expected real interest rate increases. Second, the parameter ρ is overestimated by anywhere from zero to 15 percent, depending on the expected real interest rate. Third, unlike the case of risky labor income discussed above, the ratio does not seem to change much from period to period.

The remaining figures show how this ratio changes when the values of various parameters are changed from the baseline case. For example, Figure 13 shows that when the standard deviation of *e* is lower (set equal to 0.01 in the figure), the bias in ρ is always upward, but the magnitude is small for the range of interest rates implied by the lower variance. With σ_e as low as 0.01, ρ is overstated by 8 percent at most, when the expected real interest rate is around 8 percent. Figures 14 and 15 show how the bias changes when ρ is assumed to be higher than the base case, with values of 6 and 8 respectively. Again, the bias is generally positive, and ranges from zero to 15 percent of what is implied by the Taylor expansion, depending on the expected real interest rate. As in the other cases, the bias increases with the expected real interest rate. Figure 16 shows the same graph when the AR coefficient is negative, equal to -0.1, and the standard deviation of e is 0.15, mimicking the annual return on an indexed stock portfolio. Negative autocorrelation reduces the variance in the expected returns, which appears to limit the bias relative to the base case; ρ is overestimated by 2 to 8 percent as the interest rate rises.

The results illustrate several points. First, not only is the direction of bias in ρ the opposite of what was found in the risky labor income model, the overall degree of bias is smaller. Results in the last section indicated that ρ could be underestimated by as much as 60 percent of the true value, whereas findings in this section suggest that ρ is generally overestimated by no more than 15 percent. Note that there is no reason to expect the bias to be similar in the two cases since the models are fundamentally different. Second, like the risky labor income model, the overall degree of bias is not constant, and depends on a number of economically meaningful variables such as the current interest rate and how wealthy the individual is. Third, and also like the risky labor income case, more variance in the driving process results in more bias in the estimated coefficient.

IV. Conclusion

This paper demonstrates that the use of linear approximation of dynamic Euler equations can be risky. In both of the models we consider, approximations can result in incorrect inferences about the degree of risk aversion and the importance of precautionary saving motives. Furthermore, our results help to resolve some puzzles in literature on consumption behavior. For example, the empirical literature on models of precautionary saving motives indicates implausibly small values of the degree of risk aversion (see Dynan, 1993). However, our results indicate that approximation error is likely to result in measures of the degree of risk aversion that are biased down. On the other hand, there is a large literature that examines how expected real interest rates affect consumption growth, and estimates of these models imply a very large coefficient of relative risk aversion. Our results indicate that in this case approximation error is likely to result in measures of risk aversion that are biased up.

Our results may also explain why estimates of parameters of the utility function (using micro-level data) may 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. Furthermore, the downward bias in the coefficient of relative prudence is inversely related to the degree of true prudence. We are therefore faced with the irony that the more prudent individuals are, the more we will underestimate the true degree of precautionary behavior.

The findings strongly suggest that studies which focus on assessing the importance of liquidity constraints in linearized Euler equations by splitting the sample according to wealth (or some other indicator of whether households are likely to be liquidity constrained) will lead to incorrect inferences about the degree to which the consumption growth of various subsamples is excessively sensitive to predictable components of income. This point has be made by Zeldes (1989b) and Carroll (forthcoming) with respect to omitting the *variance* of consumption growth from the linearized equation. The findings here are relevant to this body of work because they show, not only that higher order moments may be important omitted terms, but also how the bias varies with a number of economically meaningful variables. The problem is that predictable income growth will be correlated with omitted terms which are the source of linearization bias. The linearization bias itself varies with wealth and income uncertainty and prudence, the same factors that determine how willing individuals are to borrow, and how able they are to buffer transitory shocks to their income without desiring to borrow[8]. In short, the same variables that influence the degree of linearization bias also influence the probability of being liquidity constrained. This makes it very unclear whether it is actually possible to split the sample in such a way that successfully identifies groups as constrained or unconstrained on the one hand, and holds fixed the linearization bias across the subsamples on the other.

It is not clear how these problems can be handled empirically. One "solution" would be to derive explicit measures of the higher-order terms that are the source of the bias. In practice this is infeasible, simply because there are an infinite number of terms. It would be useful to explore, in future work, whether controlling for a handful of higher-order moments reduces the extent of the bias. However, this approach is unlikely to be worthwhile for the macro literature, discussed above, that uses aggregate data to examine how expected interest rate movements affect consumption growth.[9] An alternative is to use GMM estimates that do not rely on linear approximations to the Euler equation. In practice, however, other problems arise in applying the

---

[8]Ludvigson (1996) documents how variation in certain key preference and uncertainty parameters affects how often an individual will be constrained, in a model of forward looking consumption choice subject to a time-varying borrowing limit.

[9] To see the problem, consider the third term on the right-hand-side of equation (9). This term would in theory be measured as the average, over all individuals, of the variance of consumption growth, rather than the variance of aggregate consumption growth. In the presence of idiosyncratic (i.e. individual-specific) risk, the two measures will be quite different.

GMM approach, the most notable difficulty being that overidentifying restrictions are often strongly rejected, suggesting that there is something else amiss with the model (e.g., see Hansen and Singleton, 1982; Campbell, Lo, MacKinlay, 1997, chapter 8).[10] The findings in this paper show that even if the econometrician has the correct model of consumer behavior, estimates of key preference parameters may be biased if the estimating growth equation is linearized. The results therefore present an additional challenge to empirical researchers who, in practice, face a number of potentially conflicting specification issues.

[10]Campbell and Mankiw (1989) have shown that overidentifying restrictions are not rejected when one controls for predictable income growth in the *linearized* equation. This indicates that it may be necessary to account explicitly for the presence of liquidity constraints or other types of market frictions in the *nonlinearized* Euler equation when attempting to estimate parameters of the utility function. Unfortunately, this proposition is not entirely straightforward since the shadow price associated with these constraints is unobservable.

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.

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., 1993, "Intertemporal asset pricing without consumption data", *American Economic Review*, June, 487-12.

Campbell, John Y., and Andrew W. Lo, and A. Craig MacKinlay, 1997, *The econometrics of financial markets*, Princeton, NJ, Princeton University Press.

Carroll, Christopher D., 1994, "How does future income affect current consumption?", *Quarterly Journal of Economics*, February, 111-47.

Carroll, Christopher D., (forthcoming) "Buffer-Stock saving and the Permanent Income Hypothesis", *Quarterly Journal of Economics*.

Carroll, Christopher D., and Miles Kimball, 1996, "On the concavity of the consumption function", *Econometrica*, LXIV, 981-92.

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.

Epstein, L., and S. Zin, 1991, "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical investigation", *Journal of Political Economy*, 99, 263-86.

Fama, Eugene F., and Kenneth R. French, 1988, "Permanent and Temporary Components of Stock Prices", *Journal of Political Economy*, 96(2), 246-73.

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.

Hansen, Lars, and Ken Singelton, 1982, "Generalized Instrumental Variables estimation of nonlinear rational expectations models", *Econometrica*, 50, 1269-88.

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.

Ludvigson, Sydney C., 1996, "Consumption and Credit: A Model of Time-Varying Liquidity Constraints", working paper, Federal Reserve Bank of New York.

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.

Merton, Robert C., 1969, "Lifetime portfolio selection under uncertainty: the continuous time case", *Review of Economics and Statistics*, 51, 247-57.

Pischke, Jörn-Steffen, 1995, "Individual income, incomplete information, and aggregate consumption", *Econometrica*, 63(4), 805-40.

Samuelson, Paul A., 1969, "Lifetime portfolio selection by dynamic stochastic programming", *Review of Economics and Statistics,* 51, 239-46.

Weil, P., 1989, "The equity premium puzzle and the risk-free rate puzzle", *Journal of Monetary Economics*, 24, 401-421.

Wirjanto, Tony S., 1994, "Aggregate consumption behavior and liquidity constraints: the Canadian evidence", *Canadian Journal of Economics*, 28(4b), 1135-52.

Zeldes, Steve P., 1989a, "Consumption and liquidity constraints: An empirical investigation", *Journal of Political Economy*, 97, 305-46.

Zeldes, Steve P., 1989b, "Optimal consumption with stochastic income: deviations from certainty equivalence", *Quarterly Journal of Economics*, CIV, 275-98.

Table 1
Weighted averages of ratio of $(\rho^*+1)/2$ to $(\rho+1)/2$

Time period	Baseline	$\sigma = .125$	ϕ=.20	ϕ=.26 σ=.19	$\rho = 2$	$\rho = 4$	$r = .05$	$\mu = .01$
T-5	.76	.97	.68	.84	.87	.66	.75	.75
T-4	.75	.97	.68	.84	.87	.66	.76	.75
T-3	.74	.97	.68	.84	.87	.66	.76	.75
T-2	.73	.96	.68	.83	.86	.66	.77	.75
T-1	.68	.92	.65	.80	.84	.62	.74	.71

Notes: The parameters for the baseline case are: σ=.25, ρ=3, ϕ=.44, r=.03 and μ=.02. The method of computing weighted averages is described in the text.

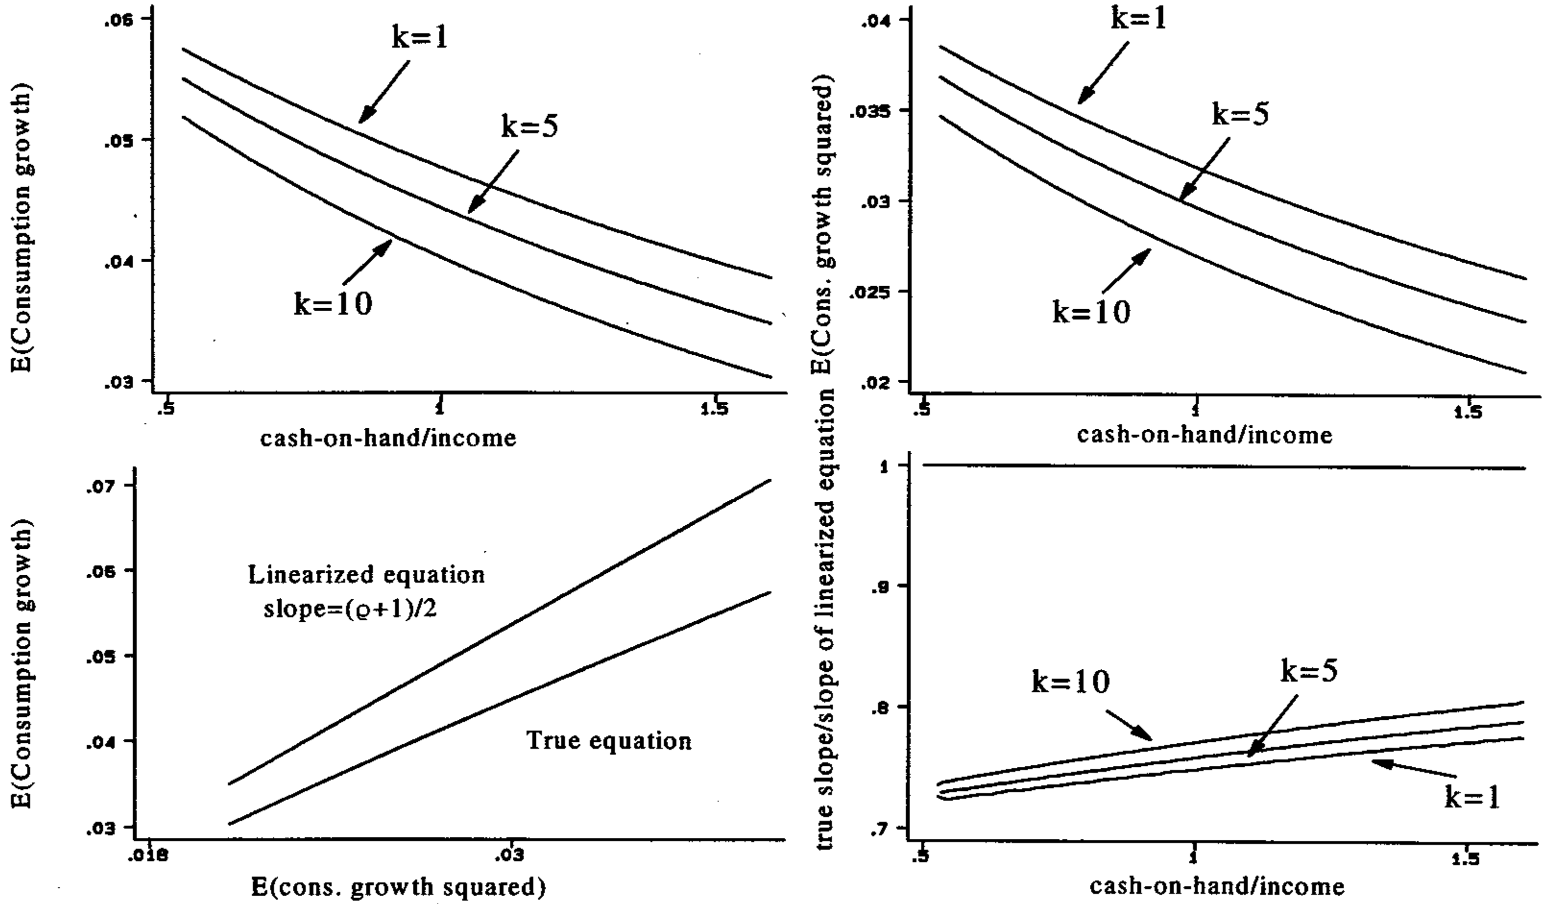

Figure 1: Baseline case, T-6

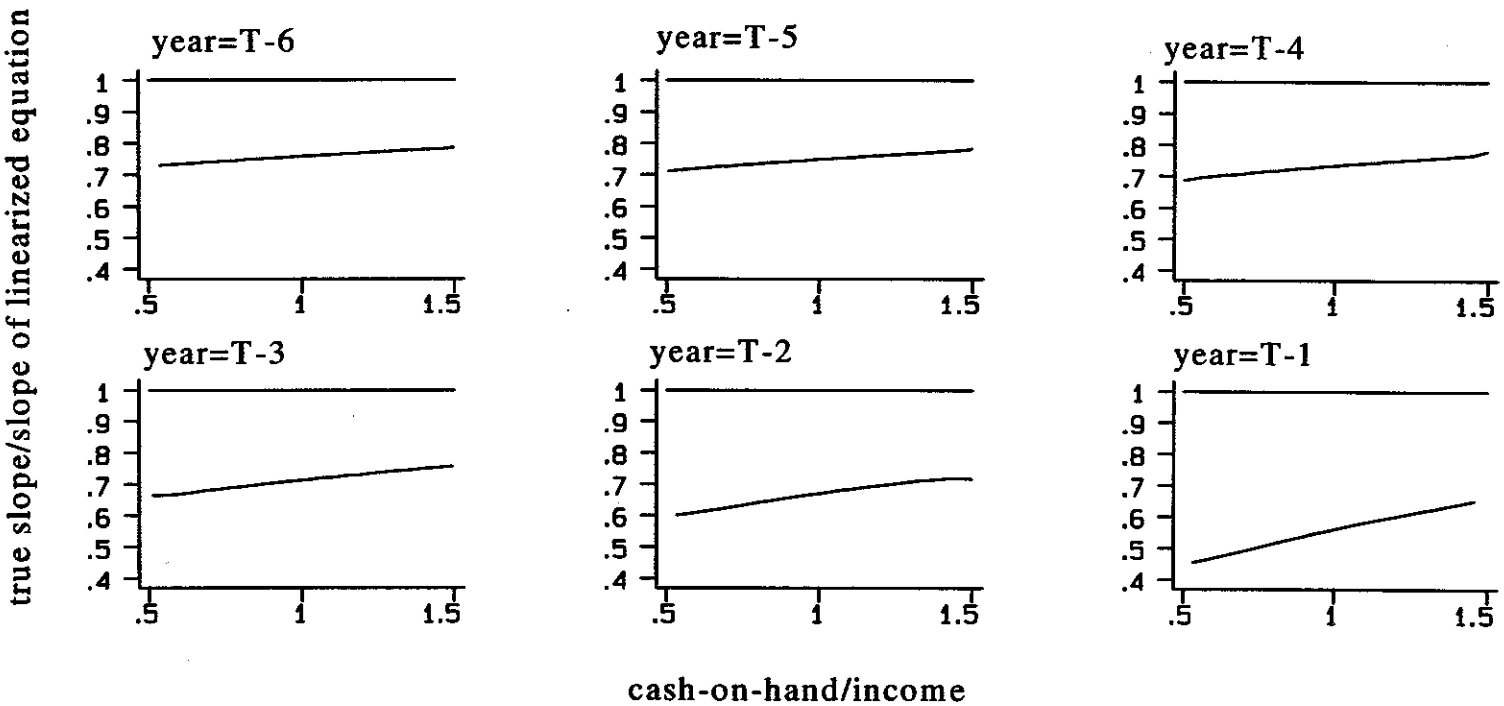

Figure 2: Baseline case

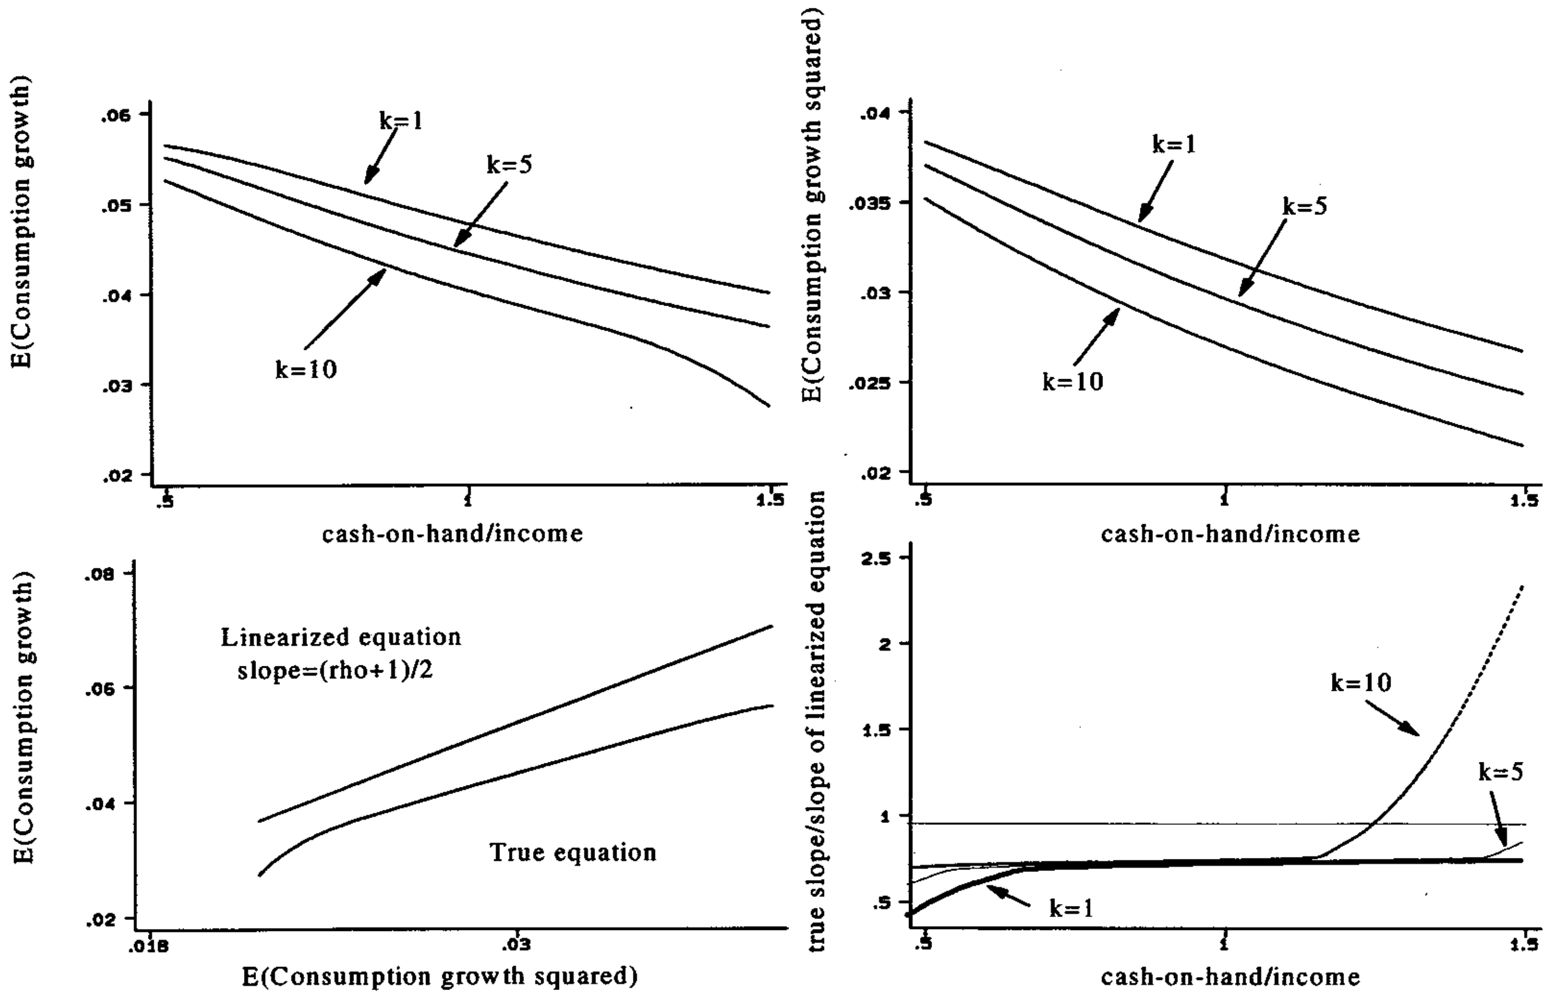

Figure 3: Baseline case, T-6, using linear extrapolation

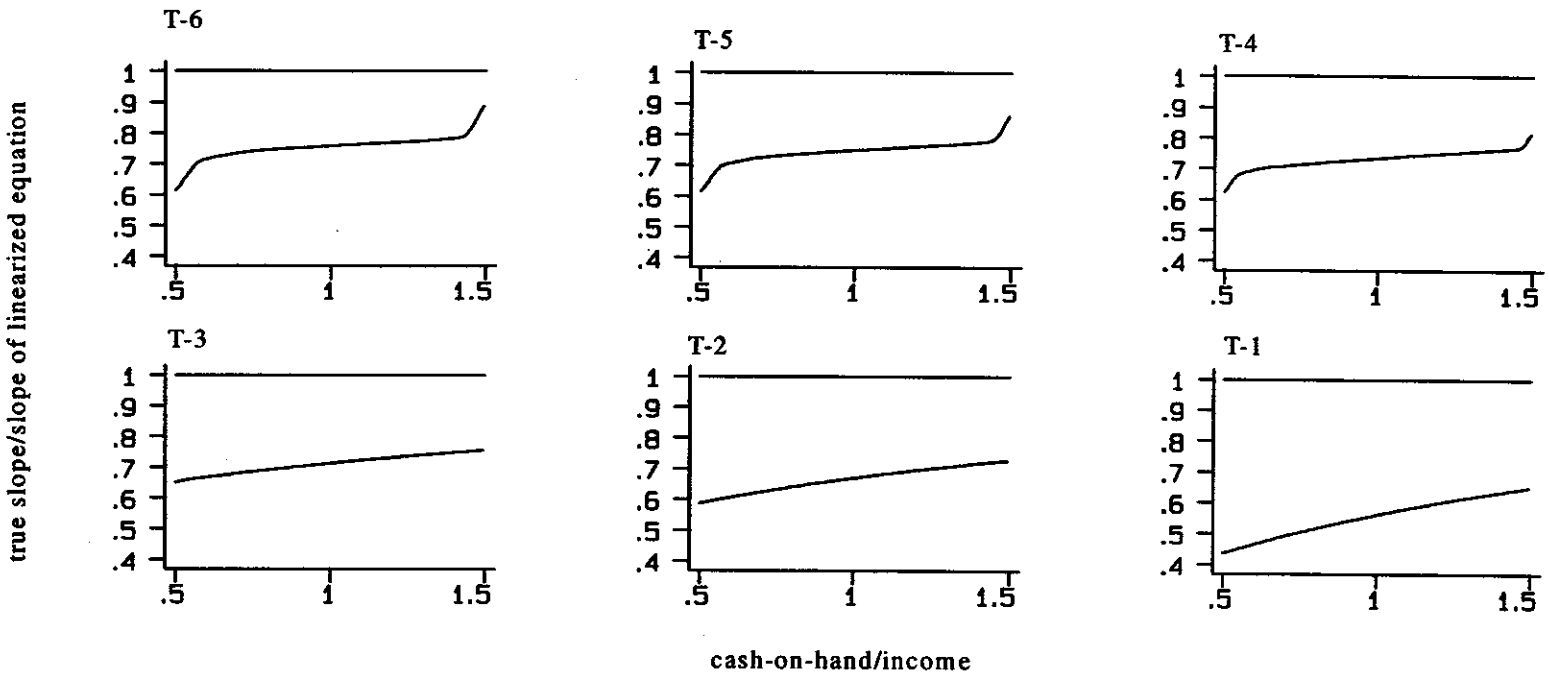

Figure 4: Baseline case, using linear extrapolation

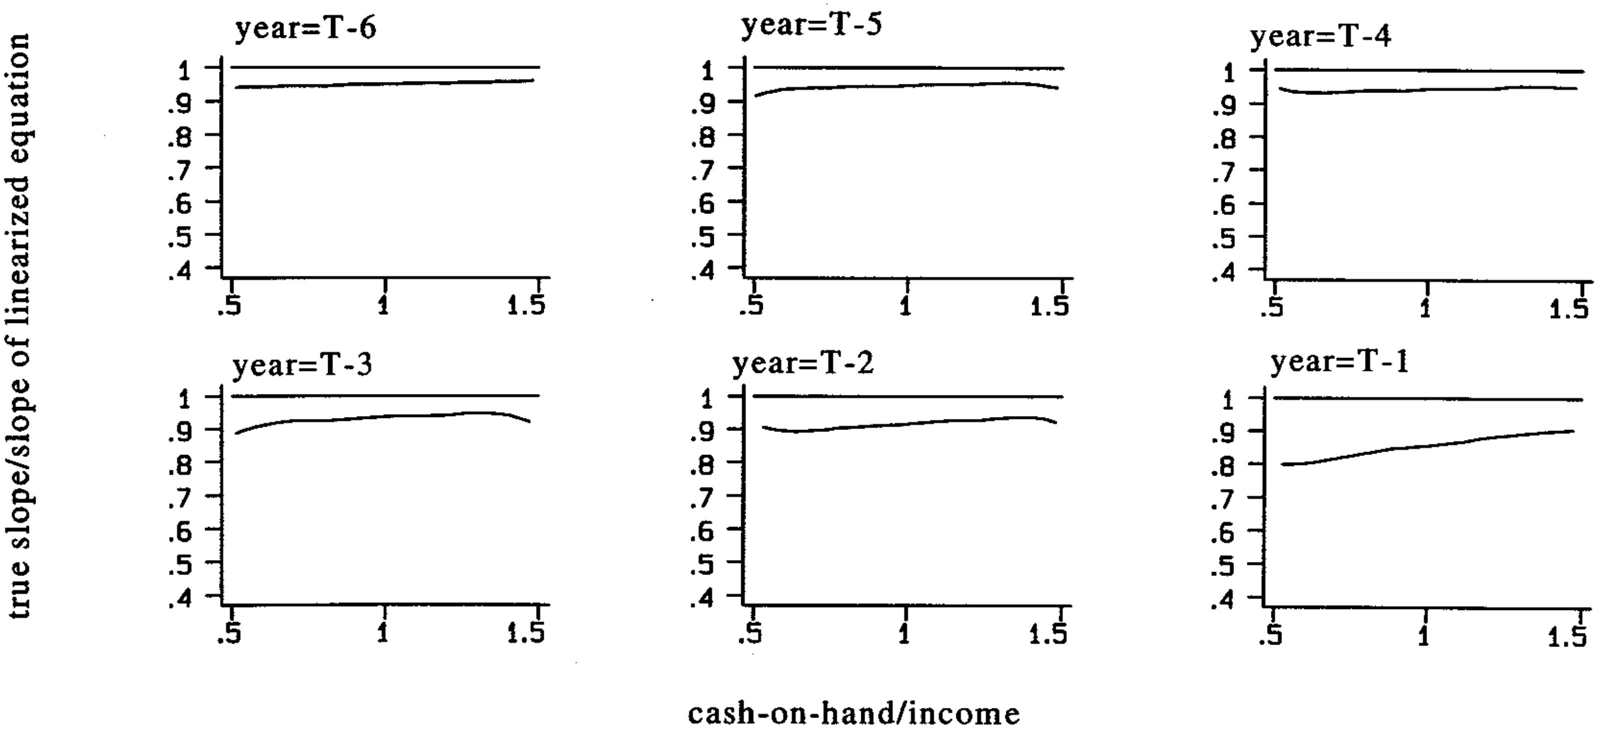

Figure 5: $\sigma = .125$

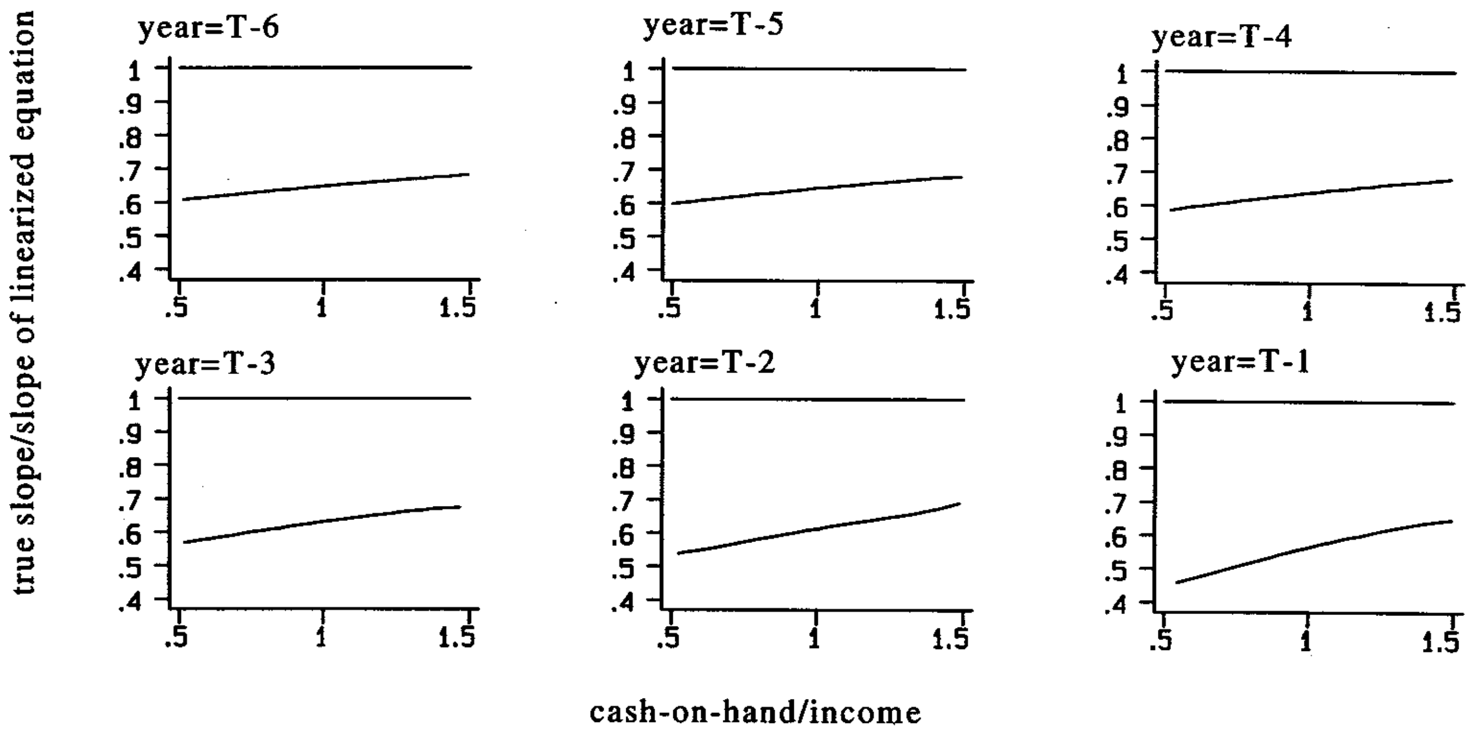

Figure 6: ϕ=.20

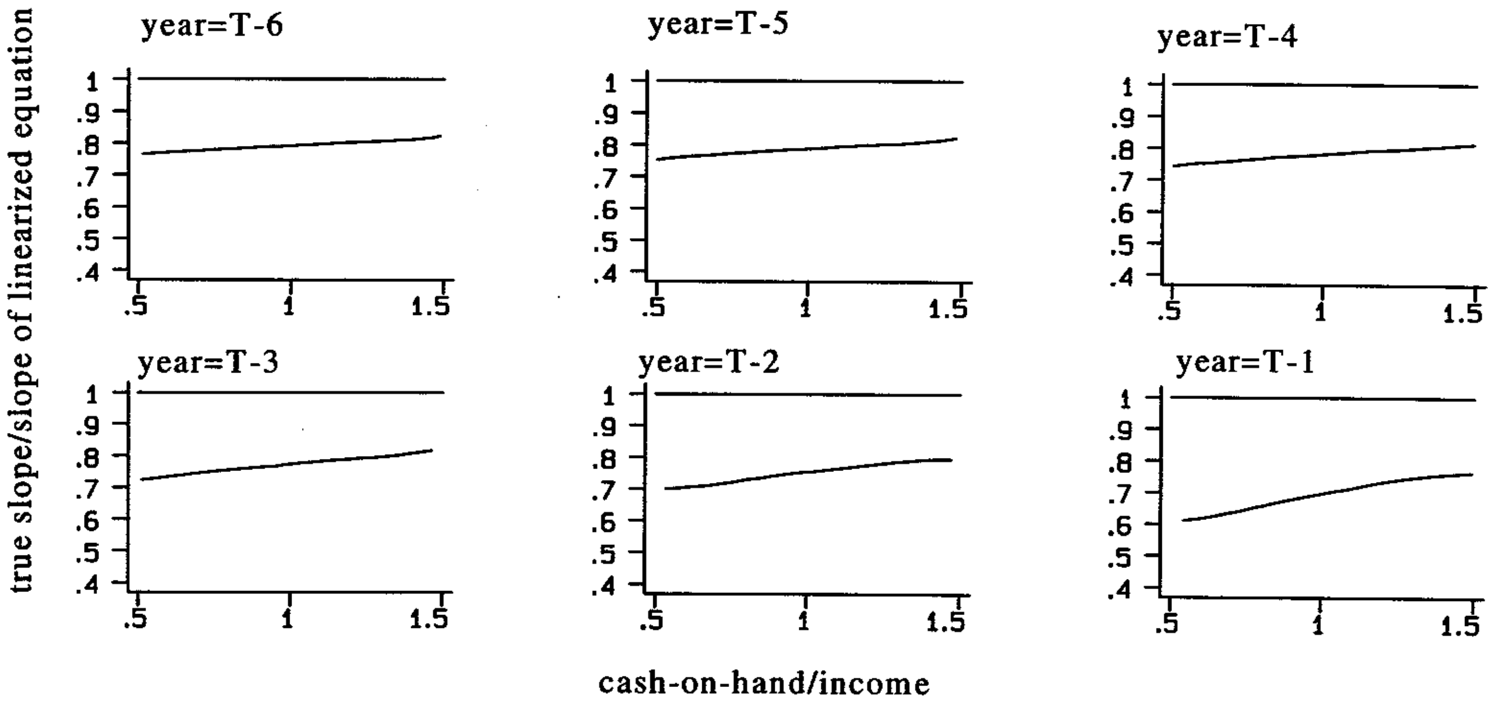

Figure 7: ϕ=.26, σ=.19

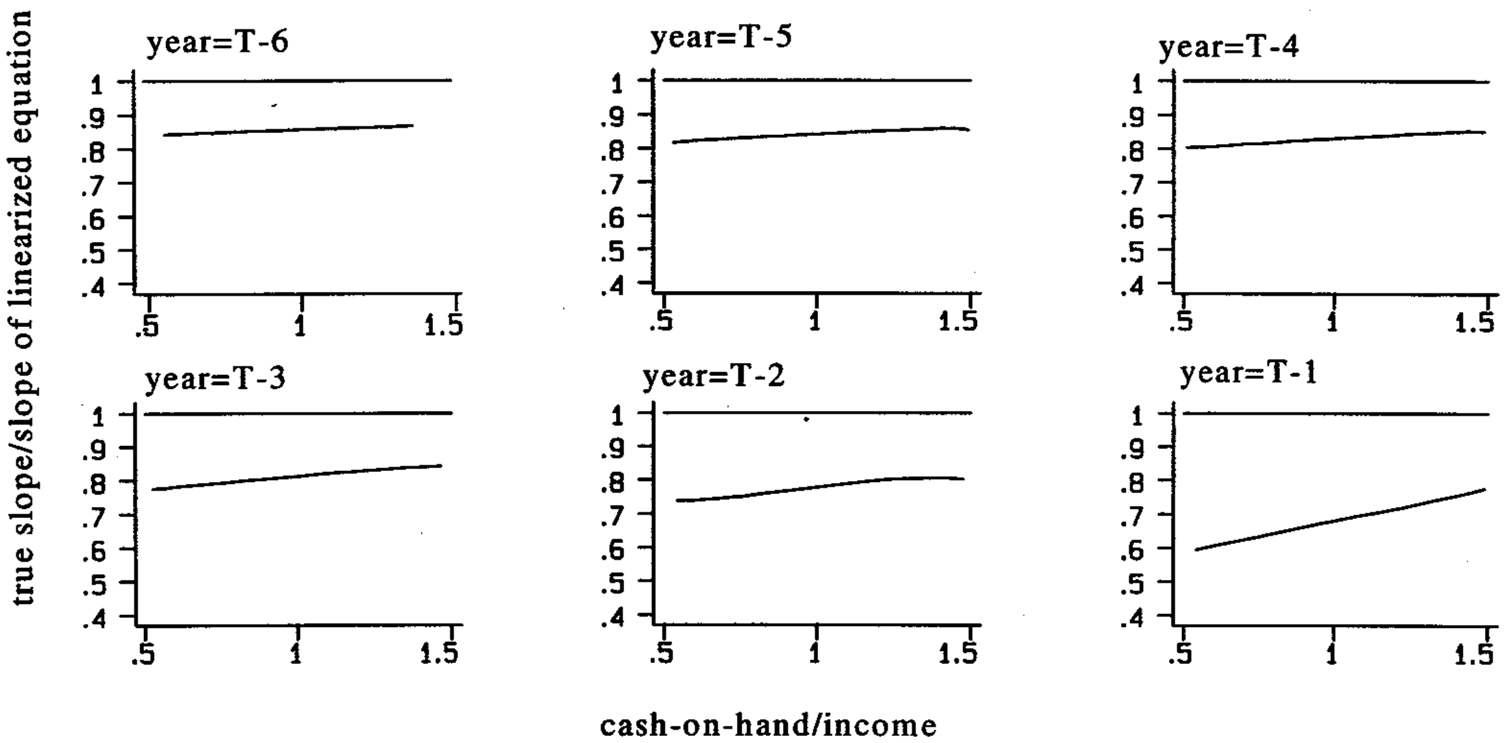

true slope/slope of linearized equation
year=T-6
year=T-5
year=T-4
year=T-3
year=T-2
year=T-1
1
.9
.8
.7
.6
.5
.4
.5
1
1.5
cash-on-hand/income

Figure 8: ρ=2

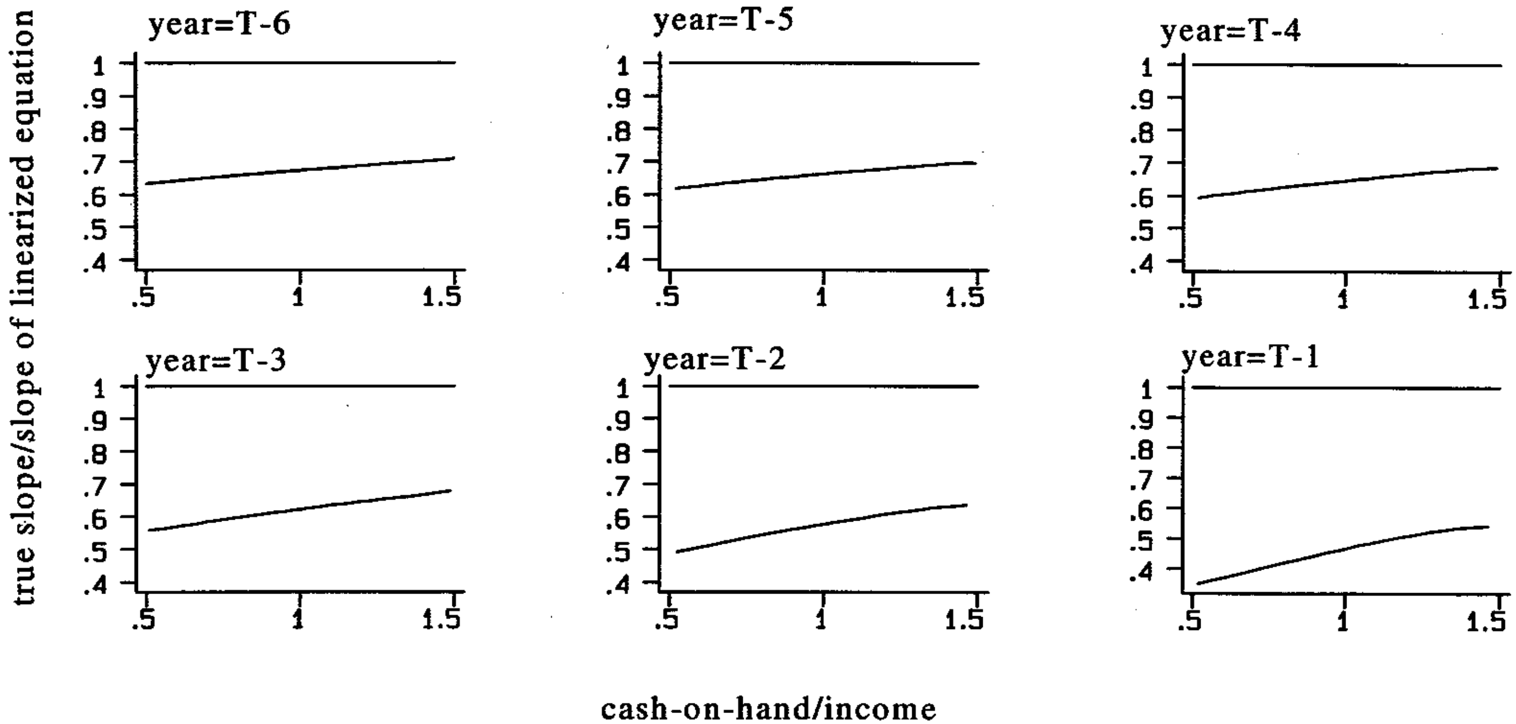

Figure 9: $\rho = 4$

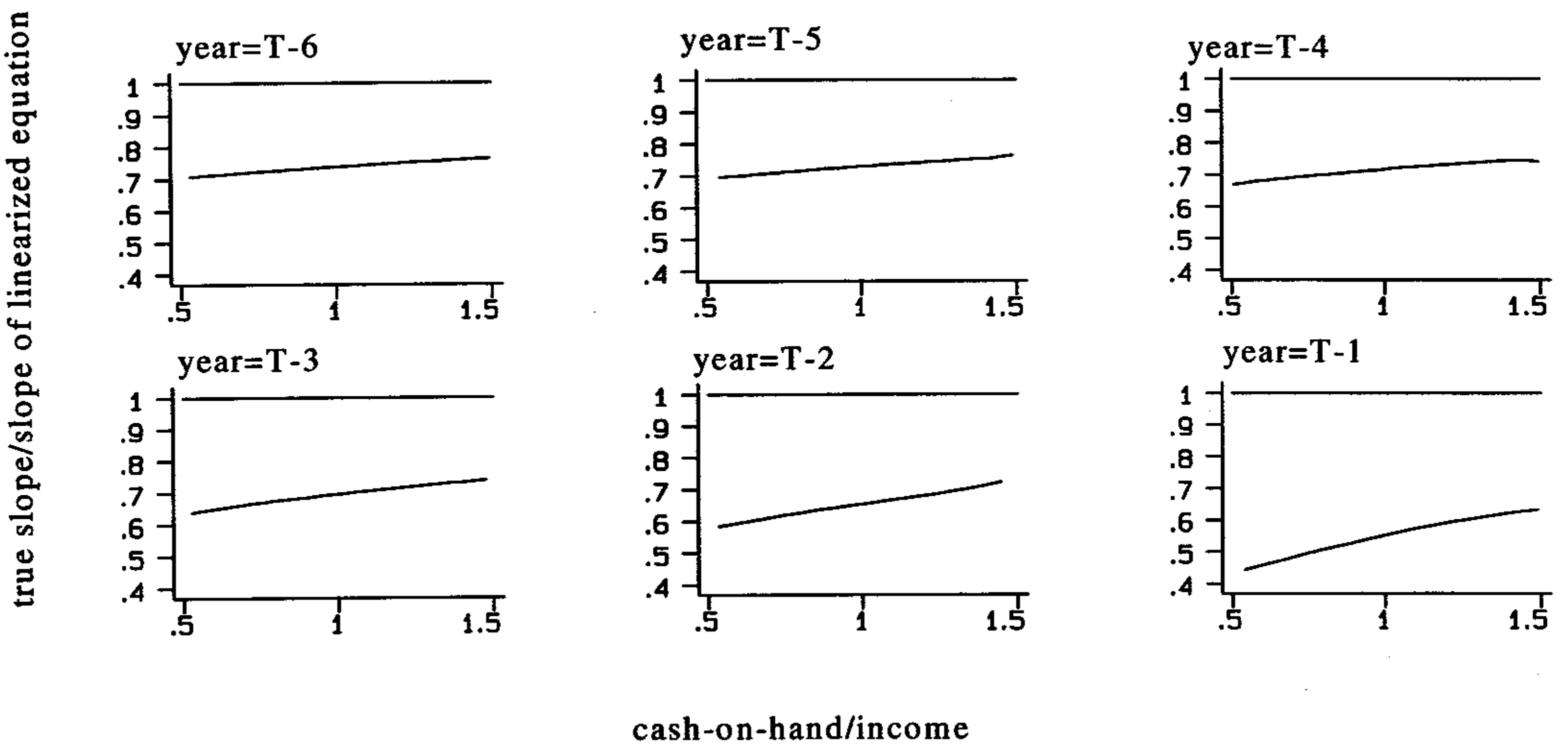

Figure 10: r=.05

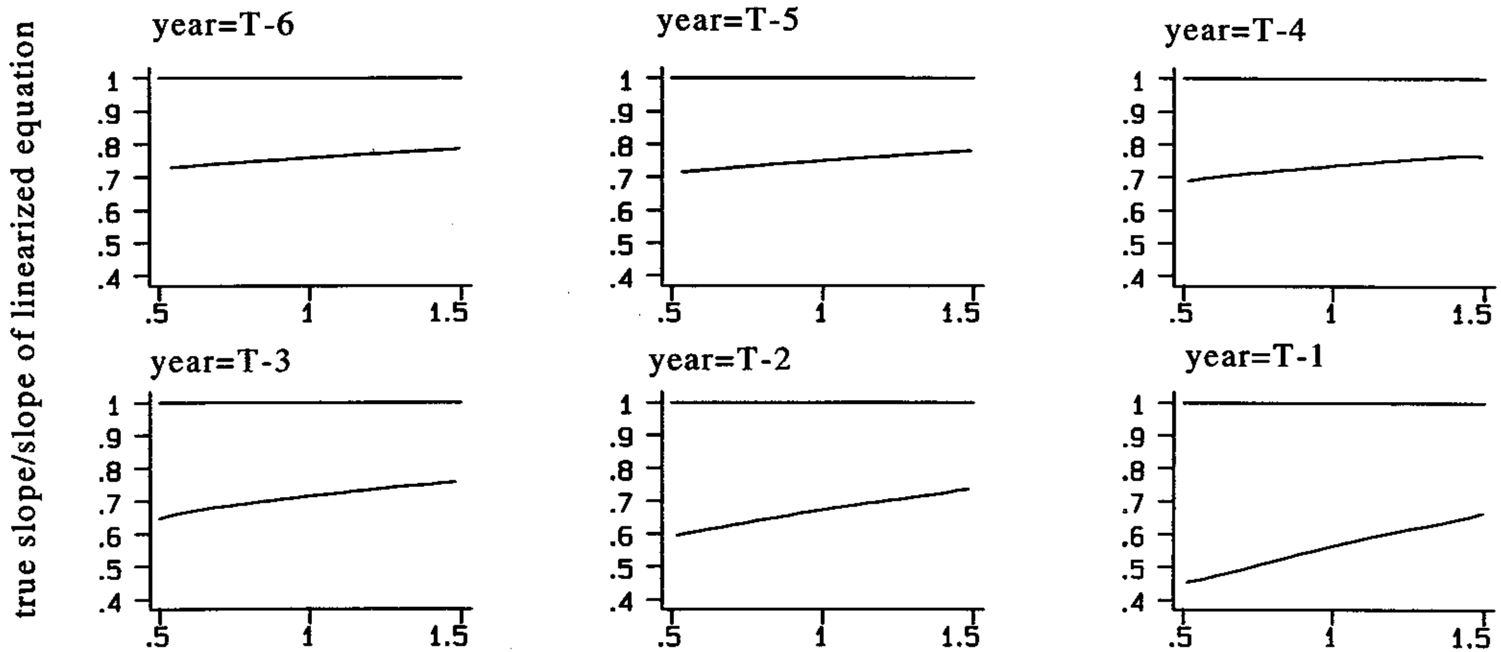

Figure 11: μ=.01

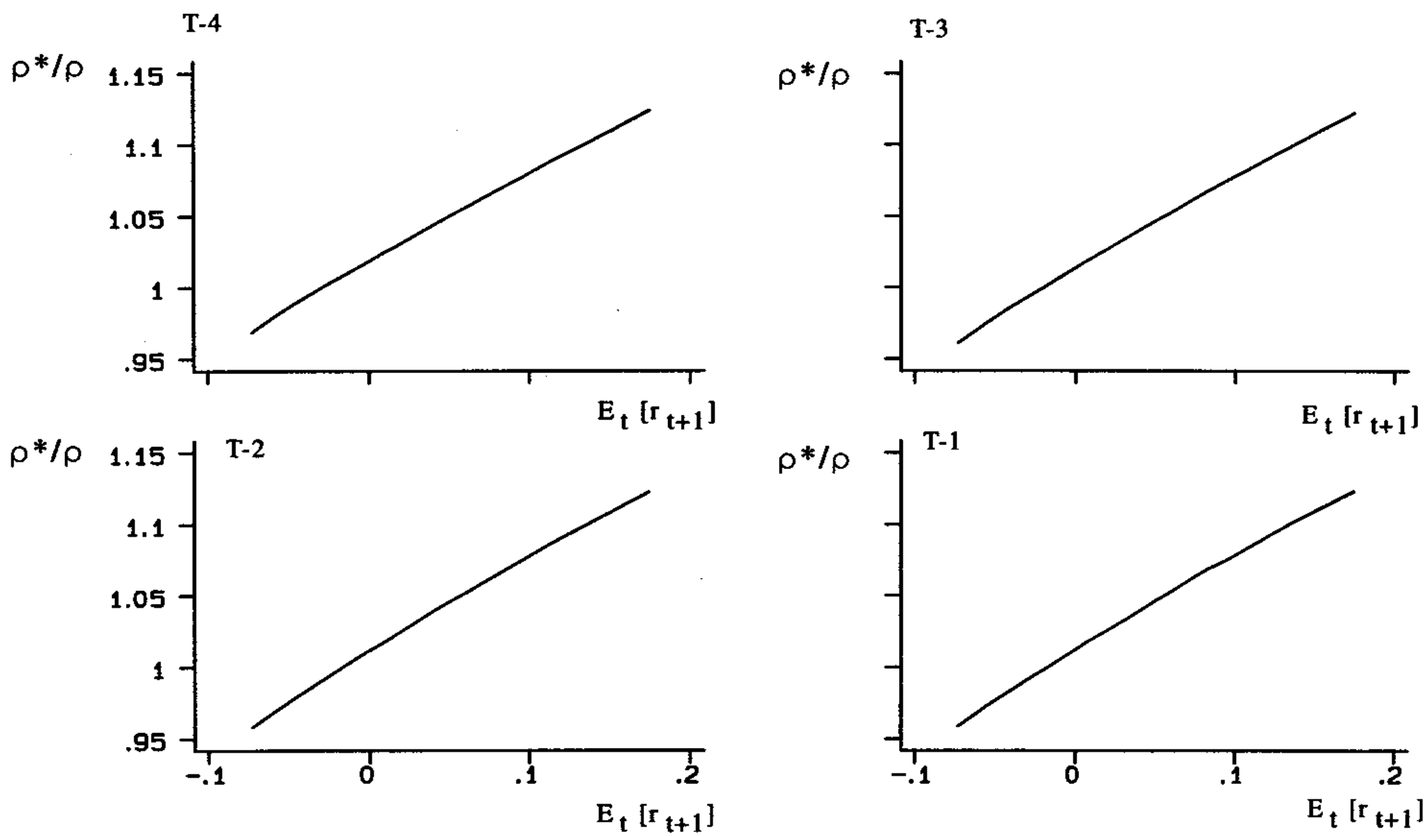

Figure 12: Risky interest rate, baseline case (α=.9, σ_ε=.05, ρ=3)

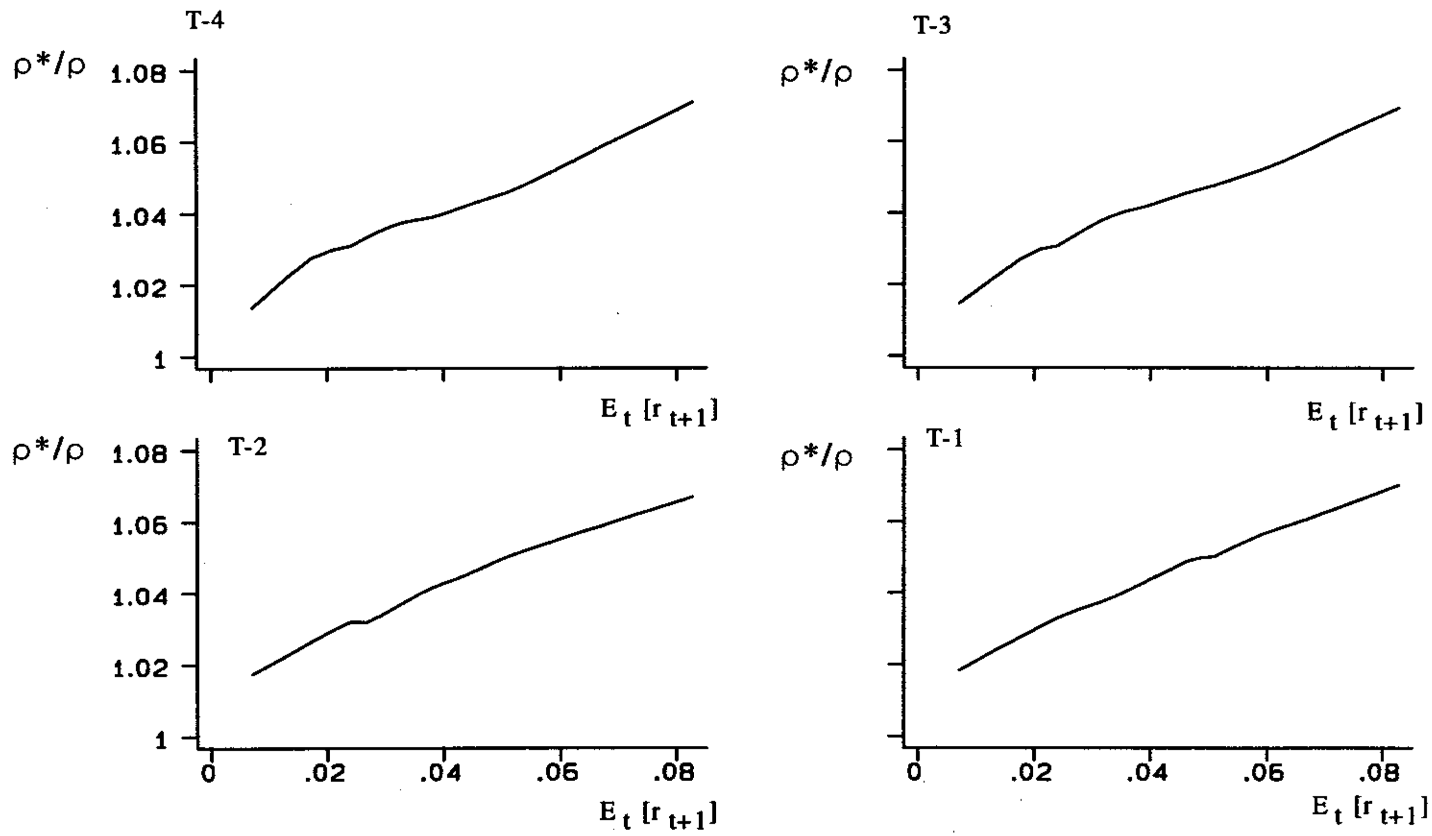

Figure 13: Risky interest rate, σ=.01

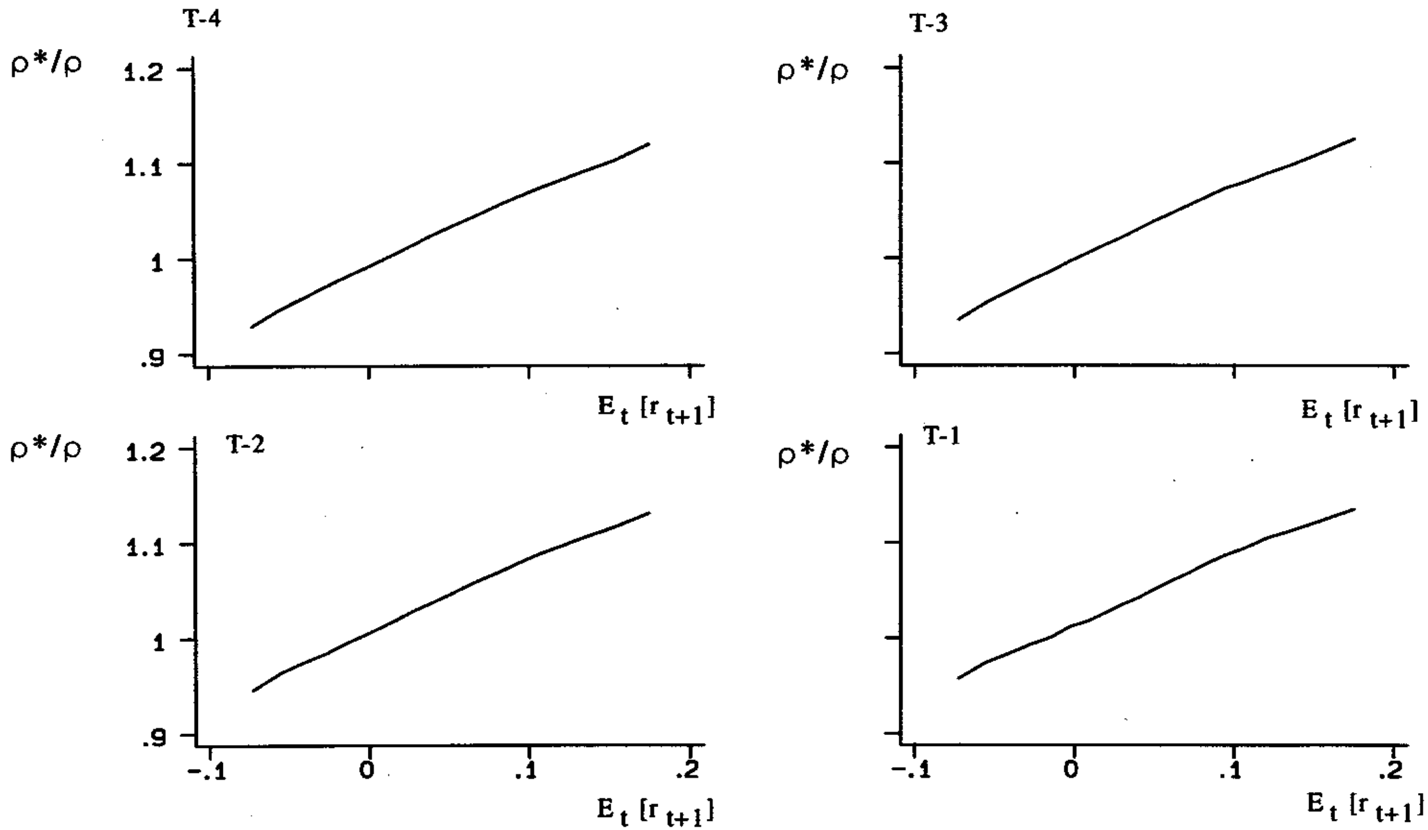

Figure 14: Risky interest rate, ρ=6

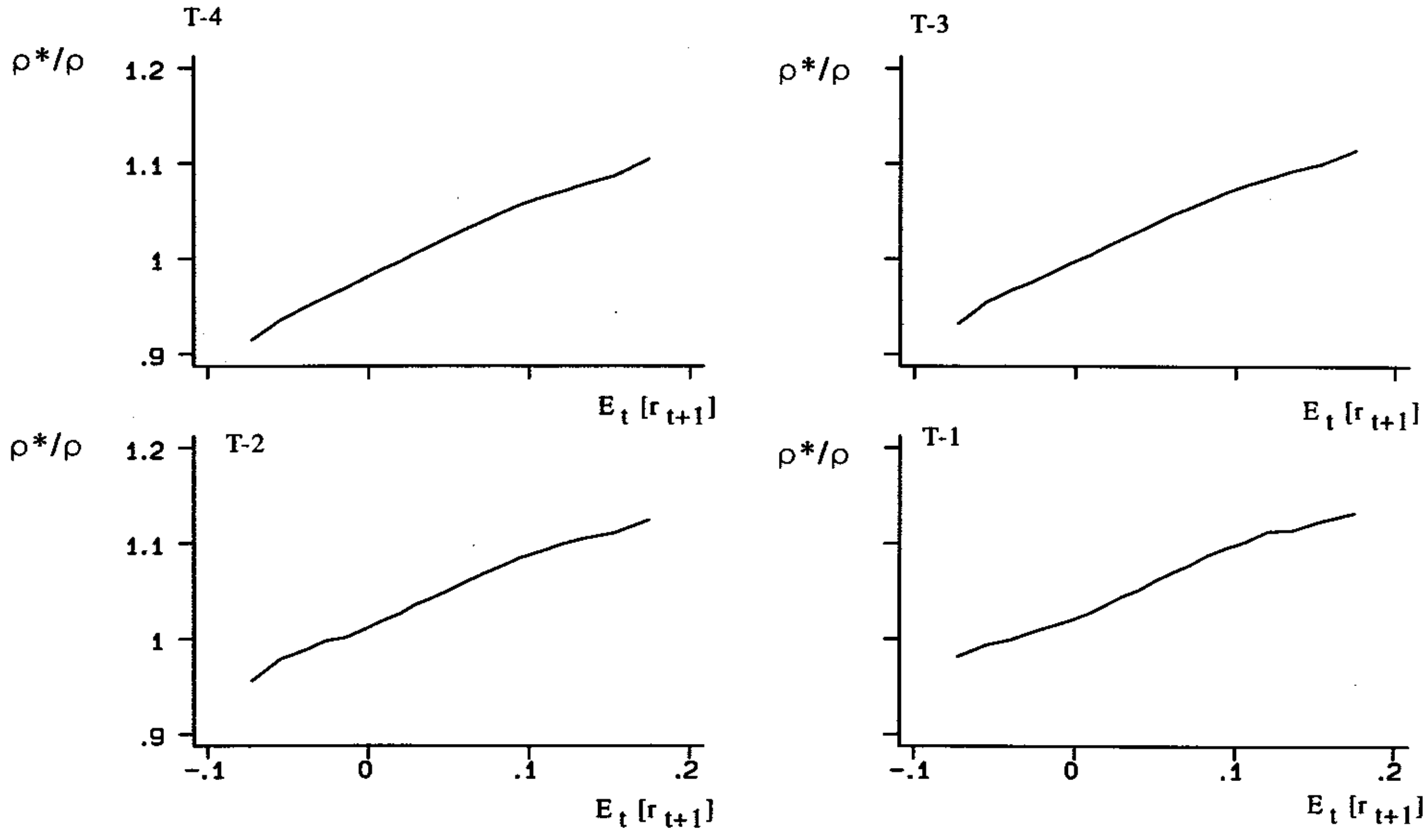

Figure 15: Risky interest rate, ρ=8

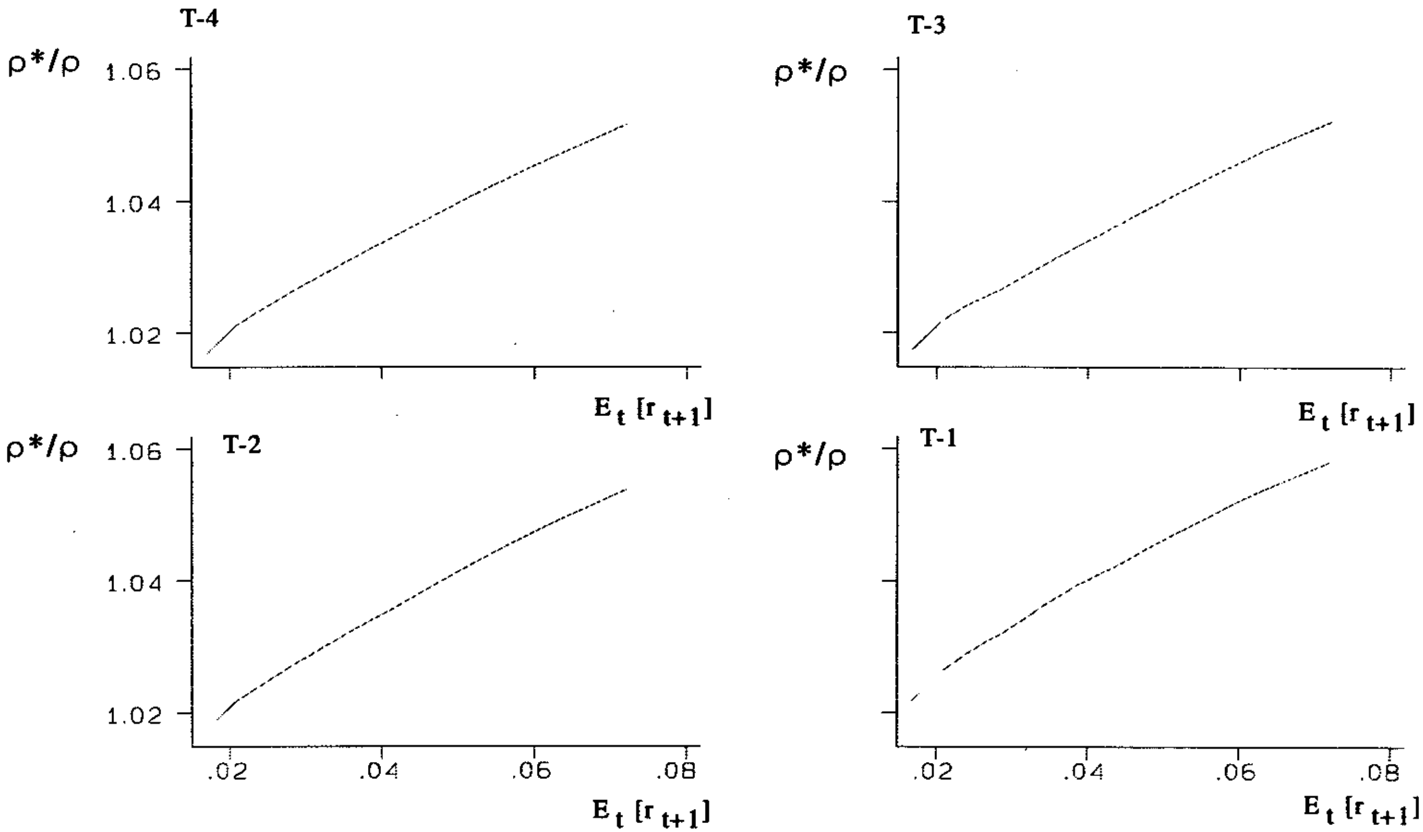

Figure 16: Risky interest rate, α=-.10, σ_e=.15

FEDERAL RESERVE BANK OF NEW YORK
RESEARCH PAPERS
1997

The following papers were written by economists at the Federal Reserve Bank of New York either alone or in collaboration with outside economists. Single copies of up to six papers are available upon request from the Public Information Department, Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045-0001 (212) 720-6134.

9701. Chakravarty, Sugato, and Asani Sarkar. "Traders' Broker Choice, Market Liquidity, and Market Structure." January 1997.

9702. Park, Sangkyun. "Option Value of Credit Lines as an Explanation of High Credit Card Rates." February 1997.

9703. Antzoulatos, Angelos. "On the Determinants and Resilience of Bond Flows to LDCs, 1990 - 1995: Evidence from Argentina, Brazil, and Mexico." February 1997.

9704. Higgins, Matthew, and Carol Osler. "Asset Market Hangovers and Economic Growth." February 1997.

9705. Chakravarty, Sugato, and Asani Sarkar. "Can Competition between Brokers Mitigate Agency Conflicts with Their Customers?" February 1997.

9706. Fleming, Michael, and Eli Remolona. "What Moves the Bond Market?" February 1997.

9707. Laubach, Thomas, and Adam Posen. "Disciplined Discretion: The German and Swiss Monetary Targeting Frameworks in Operation." March 1997.

9708. Bram, Jason, and Sydney Ludvigson. "Does Consumer Confidence Forecast Household Expenditure: A Sentiment Index Horse Race." March 1997.

9709. Demsetz, Rebecca, Marc Saidenberg, and Philip Strahan. "Agency Problems and Risk Taking at Banks." March 1997.

9710. Lopez, Jose. "Regulatory Evaluation of Value-at-Risk Models." March 1997.

9711. Cantor, Richard, Frank Packer, and Kevin Cole. "Split Ratings and the Pricing of Credit Risk." March 1997.

To obtain more information about the Bank's *Research Papers* series and other publications and papers, visit our site on the World Wide Web (http://www.ny.frb.org). From the research publications page, you can view abstracts for *Research Papers* and *Staff Reports* and order the full-length, hard copy versions of them electronically. Interested readers can also view, download, and print any edition in the *Current Issues in Economics and Finance* series, as well as articles from the *Economic Policy Review.*