

NBER WORKING PAPER SERIES

INTERTEMPORAL SUBSTITUTION IN CONSUMPTION

Robert E. Hall

Working Paper No. 720

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge MA 02138

July 1981

This research is supported by the National Science Foundation and is part of the research program in Economic Fluctuations of the National Bureau of Economic Research. All opinions expressed are mine alone. I thank John Caskey for outstanding assistance and Stephen LeRoy for helpful comments.

Intertemporal Substitution in Consumption

ABSTRACT

Does a higher real interest rate induce significant postponement of consumption? According to the theory developed here, this question can be answered by studying the relation between the rate of growth of consumption and expected real interest rates. In postwar data for the United States, expected real returns have declined over time in the stock market and for savings accounts. Over the same period, the rate of growth of consumption has been almost steady. The paper concludes that intertemporal substitution is weak, for if it were strong, the growth rate of consumption would have declined.

Robert E. Hall
Herbert Hoover Memorial Building
Stanford University
Stanford, California 94305

(415) 497-2215

Introduction

The magnitude of the substitution between present and future consumption induced by changes in the real interest rate is one of the central questions of macroeconomics. If consumers can be induced to postpone consumption by modest increases in interest rates, then (1) the IS curve is relatively flat and crowding-out is an important consideration, (2) the dead-weight loss from the taxation of interest is important, and (3) the burden of the national debt or unfunded social security is relatively unimportant, to name three of the many issues that rest on the intertemporal substitutability of consumption.

In contrast to most recent research, this paper attempts to estimate parameters of the representative individual's utility function, rather than parameters of the consumption function or savings function. As Robert Lucas (1976) has pointed out, there may not be anything that could properly be called a consumption or savings function--the relation between consumption, income, and interest rates depends on the wider macroeconomic context and may not be stable over time, even though consumers are always trying to maximize the same utility function. The techniques of this paper are more robust with respect to this kind of instability than are standard econometric models of consumption and savings.

The essential idea of the paper is the following: Consumers plan to change their consumption from one year to the next by an amount that depends on their expectations of real interest

rates. Actual movements of consumption differ from planned movements by a completely unpredictable random variable that indexes all of the information available next year that was not incorporated in the planning process the year before. If expectations of real interest rates shift, then there should be a corresponding shift in the rate of change of consumption. The magnitude of the response of consumption to a change in real interest expectations measures the intertemporal elasticity of substitution. All of this is set up in a formal econometric model where the assumptions are formalized and the estimation techniques rigorously justified.

Over the postwar period, there has been a downward shift in the expected real return from common stocks and savings accounts, the investments that presumably set the relevant real interest rate for most consumers. Over the same period, there has been only a small downward shift in the rate of growth of consumption. Consequently, all of the estimates presented in this paper of the intertemporal elasticity of substitution are small. Most of them are also quite precise, supporting the strong conclusion that the elasticity is unlikely to be much above 0.1, and may well be zero.

1. Theory of the consumer under uncertain real interest rates

The theory outlined in this section is based on the work of Douglas Breeden (1977,1979), and a number of other economists. Consumers maximize the expected value of an intertemporal utility function,

$$(1.1) \quad \dots + e^{-d(t-1)} u(c_{t-1}) + e^{-dt} u(c_t) + e^{-d(t+1)} u(c_{t+1}) + e^{-d(t+2)} u(c_{t+2}) + \dots$$

For the purposes of what follows, it is not necessary to make specific assumptions about the market setting of the maximization. At one extreme, the consumer could face a full set of markets in contingent commodities, and then the budget constraint would say that the sum of all the consumer's demands for the contingent claims valued at market prices would equal his endowment. At the other extreme, the consumer could be Robinson Crusoe, with a single risky investment in a real asset. Then the budget constraint would say that his holdings of the real asset could never be negative. For a further discussion of this point, see Sanford Grossman and Robert Shiller (1981).

In any case, one of the many choices facing the consumer is to spend a little less, x , in year $t-1$, invest x in one asset, and spend the stochastic proceeds in year t . Suppose that a unit investment in year $t-1$ returns e^{r_t} in year t . At the point of maximum expected utility, the consumer will have thought through all possibilities of this kind, and expected utility will reach

a maximum at $x=0$, so the derivative of expected utility with respect to x will be zero at that point. Only the terms in expected utility dated $t-1$ and t participate in this calculation, so the relevant first order condition can be written

$$(1.2) \quad \frac{d}{dx} E_{t-1} (e^{-d(t-1)} u(c_{t-1} - x) + e^{-dt} u(c_t + e^r x)) = 0$$

thus,

$$(1.3) \quad -e^{-d(t-1)} u'(c_{t-1} - x) + E_{t-1} e^{dt} u'(c_t + e^r x) = 0$$

At $x=0$,

$$(1.4) \quad E_{t-1} e^{rt} u'(c_t) = e^d u'(c_{t-1}) .$$

This is the precise mathematical formulation of the principle that the marginal rate of substitution should equal the ratio of the prices of present and future consumption. Under uncertainty, it is not true that the expected marginal rate of substitution should equal the expected price ratio (the discount function). Rather, the expected value of future marginal utility multiplied by the stochastic return should equal the current value of marginal utility. Note that this "reallocation" condition is the generalization of the proposition investigated in my earlier paper (Hall (1978)) that

marginal utility should be a trended random walk when real interest rates are constant over time.

Further progress in translating the reallocation condition into consequences for observed variables requires assumptions about the distributions of the random influences. A set of assumptions introduced by Breeden (1977) seems a natural approach. First, assume that the real interest rate, r_t , conditional on information available in year $t-1$, obeys the normal distribution with mean r_t^* and variance v_r . Because interest rates as they are defined in this paper can be indefinitely negative, the normal distribution is a natural assumption. Second, assume that the consumer's rule for processing new information about income and interest rates makes the distribution of marginal utility log-normal, conditional on information available last year; that is, $\log u'(c_t)$ is normal with mean m_t and variance v_c .

Because the new information arriving in year t has a bearing on both the actual return to investments maturing in t and on the consumer's long-term well being estimated in that year, the two random variables r_t and $\log u'(c_t)$ will be correlated; I will let $v_{r,c}$ stand for their covariance. Then the random variable on the left-hand side of the condition for optimality of the consumption rule, $e^{r_t} u'(c_t)$, is log-normal; its log has mean $r_t^* + m_t$ and variance $v_r + v_c + 2v_{r,c}$. From the rule that the expectation of the exponential of a normal random variable with mean m and variance v is $e^{m+v/2}$, the expectation of the left-hand side of the condition is

$$(1.5) \quad E_{t-1} e^{r_t} u'(c_t) = \exp(r_t^* + m_t + v_r/2 + v_c/2 + v_{r,c}) .$$

This should equal the right-hand side of the condition,

$$(1.6) \quad e^d u'(c_{t-1}) .$$

In logs, the condition takes the simple linear form,

$$(1.7) \quad r_t^* + m_t + v_r/2 + v_c/2 + v_{r,c} = d + \log u'(c_{t-1}) .$$

Recall that the condition is a constraint on the consumption rule. It may be useful to rearrange it so that those parts controlled by the consumer are on the left and exogenous parts are on the right:

$$(1.8) \quad m_t - \log u'(c_{t-1}) + v_c/2 + v_{r,c} = -r_t^* - v_r/2 + d .$$

If the structure of uncertainty is stable over time, v_c , $v_{r,c}$, v_r , and d can all be combined in a constant, k :

$$(1.9) \quad m_t - \log u'(c_{t-1}) = k - r_t^* .$$

Put another way, the random variable $\log(u'(c_t)/u'(c_{t-1}))$ is normal with mean $k - r_t^*$ and variance v_c .

The development of the model to this point is the following: r_t and $\log(u'(c_t)/u'(c_{t-1}))$ are bivariate log-normal with means

r_t^* and $k - r_t^*$, conditional on information available in year $t-1$. The strong testable implication of the theory is that the mean of the marginal rate of substitution is shifted only by the mean of the real interest rate. Information available in year $t-1$ is helpful in predicting the marginal rate of substitution only to the extent that it predicts the real interest rate. This testable implication is the logical extension of the one derived in my earlier paper under constancy of real interest rates. In that case, no variable known in year $t-1$ should help predict the marginal rate of substitution.

The next step in deriving testable implications is to assume a functional form for the utility function. Breeden has suggested the natural choice,

$$(1.10) \quad u(c) = (1-1/s)^{-1} c^{1-1/s} .$$

Its marginal utility is

$$(1.11) \quad u'(c) = c^{-1/s} .$$

Here s is the intertemporal elasticity of substitution and the reciprocal of the coefficient of relative risk aversion. In the next section, a more elaborate specification is introduced where separate parameters control intertemporal substitution and risk aversion. I show that the procedure developed here for the simple additive utility function actually estimates the intertemporal elasticity of substitution. The bivariate

relation between consumption and real interest rates does not reveal anything about risk aversion.

It is convenient at this point to multiply the log of the marginal rate of substitution by the elasticity of substitution, s , to put the bivariate model in the following form: r_t and $\log(c_t/c_{t-1})$ are jointly normal with means r_t^* and $k + sr_t^*$, variances v_r and v_c , and covariance $v_{r,c}$. Here I have redefined k , v_c , and $v_{r,c}$ so as to eliminate an inconvenient s .

Finally, it will be useful to write out the model in more standard econometric form, with explicit disturbances:

$$(1.12) \quad r_t = r_t^* + u_t$$

$$(1.13) \quad \log(c_t/c_{t-1}) = k + sr_t^* + v_t$$

2. Distinguishing intertemporal substitution from risk aversion

In this section, I will argue that the regression of the rate of change of consumption on the expected real interest rate reveals the intertemporal elasticity of substitution, not the coefficient of relative risk aversion. In order to infer anything about risk aversion, more than one asset must be considered. The argument proceeds by introducing a utility function where separate parameters control risk aversion and intertemporal substitution. Consumer optimization gives rise to exactly the same bivariate model of real returns and the change in consumption derived in the previous section. The risk aversion parameter is not identified econometrically from the data on the return to a single asset and the rate of change of consumption.

The utility function is the earlier function raised to a power:

$$(2.1) \quad - E_{t-1} \frac{1/s-1}{a-1} \left[\sum_{t'=t-1}^{\infty} e^{-dt'} \frac{c_{t'}^{1-1/s}}{1-1/s} \right]^{\frac{a-1}{1/s-1}}$$

I define the intertemporal elasticity of substitution as the elasticity of substitution between consumption in any pair of years under certainty; this is plainly the parameter, s . On the

other hand, I define the coefficient of relative risk aversion according to its standard atemporal definition applied to the utility function obtained by inserting a common stochastic level of consumption in every year into the function just given. This is easily seen to be the parameter a . I assume that s is less than one and a exceeds one.

Based on the same logic as before--that the reallocation of a unit of wealth from year $t-1$ to year t not change expected utility--the reallocation condition is

$$(2.2) \quad - E_{t-1} \left[\sum_{t'=t-1}^{\infty} e^{-dt'} \frac{c_{t'}^{1-1/s}}{1-1/s} \right]^{\frac{a-1/s}{1/s-1}} \cdot$$

$$(-e^{-d(t-1)} c_{t-1}^{-1/s} + e^{rt} e^{-dt} c_t^{-1/s}) = 0$$

Let

$$(2.3) \quad z_t = \left[\sum_{t'=t-1}^{\infty} e^{-dt'} \frac{c_{t'}^{1-1/s}}{1-1/s} \right]^{\frac{a-1/s}{1/s-1}}$$

Then the reallocation condition can be written as

$$(2.4) \quad E_{t-1} z_t (e^{rt} c_t^{-1/s} - e^d c_{t-1}^{-1/s}) = 0$$

Suppose, as before, that r_t and $\log c_t$ are distributed normally conditional on information available in year $t-1$. Suppose further that $\log z_t$ is distributed normally as well. The joint distribution of the three variables is multivariate normal with means r_t^* , c_t^* , and z_t^* , variances v_r , v_c , and v_z , and covariances $v_{r,c}$, $v_{r,z}$, and $v_{c,z}$. Then the random variable

$$(2.5) \quad \log z_t + r_t - \frac{1}{s} \log c_t$$

is normal with mean

$$(2.6) \quad z_t^* + r_t^* - \frac{1}{s} c_t^*$$

and variance

$$(2.7) \quad v_z + v_r + \frac{1}{s^2} v_c + 2(v_{r,z} - \frac{1}{s} v_{c,z} - \frac{1}{s} v_{r,c})$$

and the random variable

$$(2.8) \quad \log z_t + d - \frac{1}{s} \log c_{t-1}$$

is normal with mean

$$(2.9) \quad z_t^* + d - \frac{1}{s} \log c_{t-1}$$

and variance v_z . Proceed as before, evaluating the expectations

of the exponentials of these two random variables, equate them, and solve for the implied mean of the change in the log of consumption. It is

$$(2.10) \quad c_t^* - \log c_{t-1} = sr_t^* + \frac{s}{2}v_r + \frac{1}{2s}v_c - sv_{r,z} \\ + v_{z,c} + v_{r,c} + sd$$

Note that the mean, z_t^* drops out because it appears on both sides of the reallocation condition. Collect all of the constants here into a single constant, k . Then $\log(c_t/c_{t-1})$ is normal with mean $k + sr_t^*$. The complete bivariate model of the observed variables, stated in the usual econometric form, is

$$(2.11) \quad r_t = r_t^* + u_t$$

$$(2.12) \quad \log(c_t/c_{t-1}) = k + sr_t^* + v_t$$

This is precisely the same as derived for the earlier case where $a = 1/s$. The coefficient of relative risk aversion, a , does not appear in the joint distribution of the two observed variables. The coefficient of the expected real interest rate in the consumption equation is unambiguously the elasticity of intertemporal substitution.

Estimation of the risk aversion parameter would be possible in a multivariate system with the real returns to two or more assets. Then the magnitudes of the risk premiums together with the correlations of the returns with consumption would provide

estimates of a .

3. Expectations of the real interest rate

To complete the model it is necessary to relate the conditional mean of the real interest rate, r_t^* , to observed variables known to consumers at the time that they choose c_{t-1} . Recall that r_t^* is the mean of the subjective distribution for the real interest rate held by the typical consumer at the time consumption decisions are made for year $t-1$. What I will call the "conventional specification" for expectations has been employed frequently in macroeconomic models derived from rational expectations and, in particular, underlies the recent work of Lars Hansen and Kenneth Singleton (1981) on consumption. The conventional specification says that the mean of the subjective distribution is a linear combination of observed variables:

$$(3.1) \quad r_t^* = x_{t-1}b \quad ,$$

and the coefficients, b , are known in advance. Under this specification, the complete model of expectations and

consumption becomes a simple application of bivariate regression with parameter constraints across the equations.

The conventional approach to the characterization of expectations relies on the implicit assumption that the public has always known the coefficients, b , of the forecasting equation for the real interest rate. The least-squares estimate of b embodies information that was not actually available to consumers, because it comes from a regression with later data. The fitted value, $x_{t-1}b$, cannot really claim to be the mean of the subjective distribution at $t-1$ because it draws on information unavailable in year $t-1$. As a practical matter, the problem is apparent in the following way: The fitted value, $x_{t-1}b$, is much too good a predictor of actual real interest rates, especially rates derived from the stock market. The fitted values fluctuate far too much to interpret them as truly the means of the subjective distributions held by the typical consumer.

A more satisfactory alternative is a formal Bayesian characterization of the subjective distribution of the real interest rate. In this view, consumers begin the sample period with a prior distribution on the parameters of the subjective distribution of the real interest rate. As each year passes, they update their subjective distributions to form a suitable posterior distribution, which then plays the role of the subjective distribution of the prospective real interest rate. Within the framework of a model in the form

$$(3.2) \quad r_t = x_{t-1}b + u_t \quad ,$$

if the posterior distribution for b is normal (as well as the distribution of u_t), then r_t is distributed normally conditional on x_{t-1} , and the theory of consumption developed earlier continues to apply. The mean of r_t from this type of model is much better behaved than is the fitted value from a regression. To keep the Bayesian model simple (with only a tiny substantive effect), I will assume that the variances of the surprises, u_t and v_t , are known in advance. In year $t-1$, the consumer and the econometrician have a record extending from 1 to $t-1$. The consumer wants to infer the mean, b_{t-1} , of the posterior distribution of the coefficients governing the real return. If the surprise in consumption were uncorrelated with the surprise in the real return, the problem would be a straightforward one of univariate Bayesian regression. Suppose that the consumer started with a prior distribution on b with mean \bar{b} and precision matrix P (that is, the covariance matrix of the prior distribution is P^{-1}). Then the mean of the posterior distribution for b after accumulating evidence through year $t-1$ would be

$$(3.3) \quad b_{t-1} = \left(\frac{1}{v_u} X'X + P \right)^{-1} \left(\frac{1}{v_u} X'r + P\bar{b} \right)$$

Here X and r are the matrix and vector, respectively, of data available through year $t-1$.

Because the record of forecast errors in consumption conveys

additional information relevant for estimating b , the consumer and the econometrician actually face a problem in bivariate Bayesian regression. Let W be a matrix consisting of two columns, the first a constant and the second the history of values of the mean of the posterior distribution for r_t in past years. Then define

$$(3.4) \quad z = \begin{bmatrix} x & 0 \\ 0 & W \end{bmatrix}$$

and

$$(3.5) \quad y = \begin{bmatrix} r \\ \log(c/c_{-1}) \end{bmatrix},$$

the history of real interest rates and the rate of change of consumption to date. Let q be the stacked vector of surprises,

$$(3.6) \quad q = \begin{bmatrix} u \\ v \end{bmatrix}$$

and let the covariance matrix of q be $V \otimes I$. Let d be the combined vector of all parameters:

$$(3.7) \quad d = \begin{bmatrix} b \\ k \\ s \end{bmatrix}$$

Suppose that the consumer's prior on d is normal with mean \bar{d} and precision matrix P . Then the mean of the posterior distribution for the complete set of parameters is

$$(3.8) \quad (Z'(V^{-1} \otimes I)Z + P)^{-1} (Z'(V^{-1} \otimes I)y + P\bar{d})$$

Let h_{t-1} be the posterior mean for the forecasting parameters for the real interest rate. The consumer and the econometrician can then compute the mean of the subjective distribution as

$$(3.9) \quad r_t^* = x_{t-1} h_{t-1} .$$

This number then becomes data for re-estimation next year and each subsequent year.

The following intuitive summary of the econometric procedure will provide a reasonably accurate guide: For each year, run a regression using data available only up to that year to estimate the parameters, b . Include the consumer's prior beliefs held at the beginning of the period in this regression. Compute the fitted value from the regression and call it r_t^* . This process will yield a complete time series for r^* . The very last regression will also provide the econometrician's best

description of information about both the process generating real returns and the response of consumption, provided that the econometrician also holds the same prior beliefs as the consumer.

Because the econometrician usually presents just the sample evidence, not the posterior distribution incorporating prior beliefs, it is also useful to re-estimate at the end with conventional bivariate regression, but using the same series for r^* . The re-estimation gives sample evidence conditional on the validity of the characterization of the prior beliefs of consumers at the beginning of the sample. In what follows, it is important to distinguish between the prior beliefs attributed to consumers and those which the econometrician might hold. The purpose of Bayesian analysis here is to generate a reasonable series for expected real interest rates, not to impose anyone's prior beliefs about parameter values. In particular, at no time is any informative prior placed on the parameter of highest scientific interest, the intertemporal elasticity of substitution. Moreover, the more informative is the prior on the coefficients, b , the less precise is the derived information about the elasticity.

Data

Following are brief definitions of the data series used in this study:

c_t : real consumption of nondurables (not including services) in the fourth quarter of year t , from the U.S. national income and product accounts.

r_t : realized real return after taxes on a investment in the Standard and Poor's 500 stock portfolio on a random date in the fourth quarter of year $t-1$, liquidated one year later,

OR

realized real return after taxes from a savings account earning the regulated passbook interest rate,

OR

realized real return after taxes from holding a sequence of four 90-day Treasury bills over the year.

h_t : log of the S&P 500 index of share prices, deflated.

d_t : dividend yield of the S&P 500.

z_t : nominal yield of 90-day Treasury bills in the third quarter

q_t : nominal regulated passbook interest rate in the third quarter

m_t : log of the money stock (M_{1A} concept), deflated.

p_t : log of the implicit deflator for consumption of nondurables (used as deflator for all deflated variables).

y_t : log of disposable income, deflated.

Data are observed annually, but consumption is measured as the average flow over a calendar quarter. The theory as developed earlier applies to consumption measured instantaneously, separated by a time span of any length. In practice, the time span should be long enough to permit consumers to assimilate information and put consumption plans into effect; a year seems reasonable from this point of view. Hansen and Singleton (1981) make use of unpublished monthly data on consumption, which might offer some further advantages. Again, an annual spacing of observations seems most harmonious with the theory.

It is important that the variables used by consumers to predict real interest rates (h_{t-1} , d_{t-1} , z_{t-1} , q_{t-1} , m_{t-1} , p_{t-1} , and y_{t-1}) be known when consumption dated $t-1$ is determined. For quarterly series (p and y) I used data for the third quarter. They are not actually published until about three

weeks into the fourth quarter, but this does not appear to be an important problem. For m I used data for September; again, these are not published until the first week of the fourth quarter. The stock market index, h , is published essentially instantaneously, but for the results in this version of the paper I used its average value over the third quarter. The timing of the dividend yield series, d , is ambiguous. I used the value reported by Standard and Poor's for the third quarter, but this seems to involve some averaging with earlier data, and a series with higher predictive power may be available. For the Treasury bill yield, z , I used monthly data for September, though, again, instantaneous data might be slightly superior. For the nominal passbook rate on savings accounts, q , I used the value specified in the regulations prevailing at the end of September.

After-tax magnitudes were calculated using the effective marginal rate under the federal personal income tax from John Seater (1980). The full nominal amount of dividends and interest was assumed to be taxed at this effective marginal rate. Capital gains and losses were assumed to be untaxed, on the grounds that the combination of low statutory rates, taxation only at realization, and forgiveness of accrued gains at death make the effective rate close to zero.

All data for the study are listed in an appendix available from the author.

5. Results

Before plunging into formal econometric results, I think it is useful to indicate why the data virtually dictate the answer that pervades of the the results of this paper, namely that the intertemporal elasticity of substitution is small. Some simple facts about the data are apparent just by taking averages over the first half of the postwar period (1952 through 1965) and the second half (1966 through 1979):

	Real return			Average growth of consumption
	stock market	passbook savings	Treasury bills	
1952-65	11.2%	1.5	1.3	2.9
1966-79	-1.6	-2.0	-0.2	2.5

All three measures of real returns were lower in the second half of the era, yet the growth of consumption was almost unchanged. A very rough estimate of the intertemporal elasticity of substitution is the ratio of the decline in the rate of growth of consumption (0.4 percentage points) to the decline in the real return (12.8 percentage points for the stock market, 3.5 points for passbook savings, and 1.5 points for Treasury bills). These ratios are 0.031, 0.114, and 0.270, actually quite close to what emerges from formal econometric analysis.

The stock market

Some experimentation with univariate prediction equations for the real return to common stocks suggested that the following variables had substantial value in predicting the return: the dividend yield (d_{t-1} and d_{t-2}), the change in the stock price index ($h_{t-1} - h_{t-2}$), the rate of inflation ($p_{t-1} - p_{t-2}$), and the rate of growth of real income ($y_{t-1} - y_{t-2}$). I will start by discussing the results of estimating the conventional specification for the expected real return, though these results will prove defective. Applying bivariate regression to the equations for the real return and the rate of growth of consumption gives:

$$r_t = -0.06 - 4.4 d_{t-1} + 10.4 d_{t-2} - 2.9 (p_{t-1} - p_{t-2}) - \\ (0.09) (7.3) (7.1) (1.0) \\ -0.5 (y_{t-1} - y_{t-2}) - 0.87 (h_{t-1} - h_{t-2}) \\ (0.9) (0.37)$$

$$\log(c_t/c_{t-1}) = 0.028 - 0.038 r_t^* \\ (0.005) (0.043)$$

(standard errors are in parentheses)

In the second equation, r_t^* stands for the analytical expression on the right-hand side of the first equation, not the numerical values. Actual estimation was by multivariate least squares (minimization of the determinant of the residual covariance matrix). The standard errors of the residuals in the two equations are 0.153 and 0.019 respectively.

Taken at face value, these results say, first, that there are important shifts in the expected real returns in the stock market associated with variables known in advance. The hypothesis that all the coefficients on the right-hand side of the real return equation except the constant are zero is overwhelmingly rejected. Although it is true that this equation is the result of an informal specification search, every candidate in the search revealed an important predictable element in real returns. Every equation explained at least half of the variance of r_t . All agreed that expected real returns in the stock market declined over the period, a finding that confirms results reported by Eugene Fama (1980).

Second, the intertemporal elasticity of substitution, s , (the coefficient of r_t^* in the consumption equation) is estimated quite precisely to be small. The 95 percent confidence interval includes values of s only up to about 0.05.

Table 1 gives the actual and fitted values for the real return and the rate of growth of consumption, and shows clearly the problem with the conventional econometric model of expectations when applied to a variable like the realized one-year return to common stocks. Although the equation for the real rate includes only five variables, its fitted values manage to pick up an astonishing amount of the actual variability of the left-hand variable. The equation correctly signals the stock market breaks of 1962, 1966, 1969, 1973-74, and 1977. An investor who had access to this equation throughout the period could have

Table 1. Actual and fitted values for the real return to the stock market and the rate of growth of consumption

Year	Real return		Consumption	
	Actual	Fitted	Actual	Fitted
1953	-1.4	26.9	0.0	1.8
1954	40.5	30.4	2.9	1.6
1955	26.7	7.7	5.3	2.5
1956	3.0	-4.4	1.4	3.0
1957	-14.7	5.6	1.6	2.6
1958	33.7	16.4	2.5	2.2
1959	9.1	3.3	3.0	2.7
1960	-2.2	1.1	0.6	2.8
1961	23.1	12.1	3.2	2.3
1962	-11.2	-3.5	2.8	2.9
1963	19.8	19.0	1.5	2.1
1964	12.3	-6.8	5.4	3.1
1965	8.7	-6.5	6.2	3.0
1966	-14.9	-1.9	1.6	2.9
1967	18.7	13.5	1.9	2.3
1968	5.3	-6.9	4.6	3.1
1969	-14.3	-2.7	1.7	2.9
1970	-1.1	7.0	2.8	2.5
1971	9.7	10.8	1.1	2.4
1972	12.9	-2.5	5.8	2.9
1973	-26.9	-4.4	0.5	3.0
1974	-44.3	-7.0	-2.2	3.1
1975	25.3	20.8	2.7	2.0
1976	17.9	-6.3	5.7	3.0
1977	-13.0	-3.9	3.7	2.9
1978	-2.4	6.9	3.4	2.5
1979	4.5	0.0	0.9	2.8

Note: Real return is annual percent return, after taxes. Rate of change of consumption is annual percent change.

earned an outrageous return from a suitably leveraged position based on the equation. But this is only a variant of the trivial observation that perfect foresight will make any investor infinitely rich. The rule itself embodies a lot of foresight.

The expected real return inferred from the conventional specification of expectations is excessively and implausibly volatile. Armed only with the evidence actually available in each year, nobody would predict as wide fluctuations as appear in the fitted values for the real return in Table 1. Correspondingly, the finding of a small coefficient when the fitted value is the right-hand variable in the consumption equation is no surprise. By the standard argument of errors in variables, a noisy right-hand variable receives a coefficient that is biased toward zero. Nonetheless, a good deal of investigation suggests that the true coefficient of the expected real return in the consumption equation is small, even though the standard econometric technique very clearly uses an estimate of the expected real return that is badly contaminated. In the first place, the simple manipulation reported at the beginning of this section is quite robust, though inefficient. In the standard econometric framework, it amounts to the use of a single time dummy as the only predictor in the equation for the real return. Because of its simplicity, it is much less likely to introduce excess variation into the predicted real return.

The Bayesian characterization of the subjective distribution of the real return is the most promising way to enforce

reasonable behavior on the key variable, r^* . Relative to the conventional specification, it has two favorable characteristics. First, it does not attribute perfect foresight to consumers. They are viewed as forming the subjective mean, r^* , purely from information available at the time. Second, it provides a way to make consumers mildly skeptical of strong but largely untested relations between observed variables and predicted real returns. Consumers are viewed as thinking that large coefficients in b are unlikely. As it happens, estimates of s derived from the more reasonable series for r^* emerging from the Bayesian specification confirm the finding of a very small value of s .

I assume that the public believed that the expected real return in 1953 was five percent, and that this value was unaffected by any variable known in advance. In terms of the parameter vector, b , the mean of the public's prior distribution is $(.05, 0, 0, 0, 0, 0)$. I characterize the precision of their beliefs in terms of a diagonal precision matrix, with diagonal elements of the form,

$$(5.1) \quad \frac{1}{v_u} (1, 100, 100, 100, 100, 100) p^2 \quad .$$

The overall precision is controlled by the parameter p --high values of p indicate profound skepticism about large values of the coefficients, b . I note again that no informative prior is placed on the elasticity of intertemporal substitution, s .

For simplicity, I assume that the residual covariance matrix

was known to the public from the start; I take it to be

$$V \begin{bmatrix} u_t \\ v_t \end{bmatrix} = \begin{bmatrix} .0243 & .00248 \\ .00248 & .000384 \end{bmatrix}$$

The procedure, then, is to compute the bivariate formula for the posterior mean of the coefficients, b , for each year, and then to compute the mean of the subjective distribution for the real return, r^* . The series for r^* is the only thing that is saved from the computations for each year. Then the bivariate system is estimated once again without any informative prior distribution on b , treating the series for r^* as data.

Table 2 presents the series for r^* obtained by this technique for three values of the precision parameter, p : 3, 1, and .01. Of these, the most reasonable seems to be the column for $p=1$. It is very much smoother than the predicted value from the least squares results in Table 1. Even though the public is viewed as thinking that 5 percent is the likely return as of 1953, in 1954 they have been persuaded by unfavorable experience to lower their expected return to 1.8 percent, but then very favorable returns raise the posterior mean above 10 percent through 1957. From 1959 through 1966, expected real returns remain between 9 and 11 percent. Then the subjective mean declines modestly until the debacle of 1973-74, which persuades the public that expected returns have dropped to 4 or 5 percent.

Table 2. Mean of the subjective distribution for the real return in the stock market for alternative prior precisions

Year	Actual	Mean of the posterior for alternative values of the prior precision		
		p=3	p=1	p=.01
1953	-1.4	5.0	5.0	5.0
1954	40.5	4.3	1.8	-1.4
1955	26.7	7.4	14.7	-23.0
1956	3.0	9.8	17.7	22.9
1957	-14.7	6.6	12.0	24.7
1958	33.7	5.9	8.7	14.5
1959	9.1	6.7	11.5	10.1
1960	-2.2	6.8	11.3	12.3
1961	23.1	6.7	10.0	19.2
1962	-11.2	6.9	11.1	5.7
1963	19.8	7.1	9.4	26.8
1964	12.3	6.1	10.1	0.9
1965	8.7	6.4	10.3	8.0
1966	-14.9	7.0	10.2	11.6
1967	18.7	7.9	8.7	24.1
1968	5.3	5.4	9.2	-0.4
1969	-14.3	6.9	9.0	9.5
1970	-1.1	7.9	7.9	16.6
1971	9.7	7.5	7.5	14.3
1972	12.9	5.4	7.5	1.3
1973	-26.9	5.8	7.8	2.6
1974	-44.3	6.9	6.4	-3.2
1975	25.3	10.6	4.4	23.3
1976	17.9	2.9	5.1	-9.7
1977	-13.0	3.4	5.5	-1.9
1978	-2.4	6.5	5.0	9.6
1979	4.5	5.0	4.7	1.3

The results for the consumption equation, using the r^* from Table 2, with the precision parameter, p , equal to 1, are

$$\log(c_t/c_{t-1}) = .033 - .075 r_t^* \\ (.008) \quad (.083)$$

The equation for the real return was estimated jointly, but only because of the covariance of its residuals with those of the consumption equation, so I will not trouble the reader with the parameter estimates for the real return equation. The switch to a more reasonable series for r^* only strengthens the conclusion that s is very small. Note that the estimate is quite precise. I should emphasize that no informative prior has been placed on s , only on the parameters of r^* .

The prior distribution in this analysis is not a statement about the investigator's beliefs, as in the usual application of Bayesian analysis. Rather, it summarizes what the public believed in 1953 and so presumably is related to sample evidence from earlier years. The conclusion about the low value of the intertemporal elasticity of substitution is not sensitive to the precision of the prior. With the precision increased by a factor of 3 the coefficient on r_t^* in the consumption equation becomes slightly, but not significantly, negative. With the precision decreased by a factor of 100, the coefficient is 0.044 with a standard error of 0.051. The more informative is the prior, the smoother is the r_t^* series and the higher is the

standard error of the elasticity, s . But all results agree on the low value of s .

Savings accounts

A surprisingly large volume of household wealth is held in the form of savings accounts, so it is relevant to examine the relation between their real return to and the rate of change of consumption. Recall that the basic relation derived at the beginning of the paper applies to each asset when consumers face numerous alternative means for holding wealth. For savings accounts much the same conclusion emerges as for stocks: By any reasonable measure, anticipated real returns have declined substantially over the past thirty years, while the rate of growth of consumption has remained almost constant. These two facts are consistent only with a low elasticity of intertemporal substitution.

Because the nominal return to savings accounts is tightly regulated and changes infrequently by small amounts, the main problem in predicted real returns is predicting inflation. Lagged nominal variables, particularly the money stock and lagged inflation, might seem logical candidates for predicting

the real return, and this indeed turns out to be the case. The lagged change in the nominal value of common stocks also emerged as a useful predictor over the whole sample. Because the real return to savings accounts fluctuates relatively little, the problem of wild regression coefficients and implausibly good predictions hardly arises in this case. Table 3 shows the actual and expected real return for the same prior distribution used for the stock market, but with the precision parameter, p , set to .001. The equation relating the change of consumption to the expected real interest rate is

$$\log(c_t/c_{t-1}) = 0.026 + 0.039 r_t^* \\ (0.004) \quad (0.175)$$

Again, the estimated value of the intertemporal elasticity of substitution, s , is close to zero and is reasonably precise. Because there has been rather less variation in expected real returns to savings accounts, the standard error of the estimate of s is considerably larger, but still, the confidence probability that s exceeds 0.2 is only about 15 percent.

Table 3. Actual and expected real returns for savings accounts

Year	Actual real return	Expected real return
1953	2.8	2.0
1954	1.9	2.8
1955	2.4	5.2
1956	-0.6	1.9
1957	-0.4	-1.9
1958	1.4	-0.3
1959	1.0	1.7
1960	0.6	1.3
1961	2.5	0.4
1962	1.8	0.4
1963	2.0	1.8
1964	2.4	1.1
1965	1.0	-0.2
1966	-0.5	-1.8
1967	1.2	-0.4
1968	-1.2	-2.2
1969	-1.8	-0.8
1970	-0.3	-1.1
1971	0.8	-1.0
1972	-0.1	-1.5
1973	-6.7	-2.1
1974	-9.1	-0.9
1975	-1.3	-0.1
1976	1.5	-4.1
1977	-0.7	-0.8
1978	-3.6	0.0
1979	-7.1	-2.0

Treasury bills

Direct household ownership of short-term marketable instruments like Treasury bills has been common since the mid-1960s. Again, the relation between their expected real return and the rate of growth of consumption should reveal something about the elasticity of intertemporal substitution. However, as Eugene Fama (1975) pointed out, the expected real return to Treasury bills has been close to a constant. Unlike other forms of consumer assets, there has not been a pronounced decline in the real earnings of Treasury bills. Consequently, the estimate of s derived by applying the techniques of this paper is highly imprecise. Some predictive power was found for the money stock, lagged one and two years, the lagged rate of inflation, the lagged nominal return on Treasury bills, and the lagged rate of growth of real income. With a prior mean of 1 percent in 1953 and the same precision matrix as before, with $p=.01$, the estimated relation is

$$\log(c_t/c_{t-1}) = 0.023 + 0.59 r_t^* \\ (0.005) \quad (0.54)$$

These results do not contradict the earlier findings of low values of s , but they do not support them either. The evidence from Treasury bills simply does not shed any light on the issue

of intertemporal substitutibility.

Conclusions

One cannot emerge from this study of the evidence thinking that consumption of nondurables is a major source of intertemporal substitution and therefore part of the explanation of the ups and downs of real output. This is exactly the opposite of the conclusion I reached in closely related work on intertemporal substitution in labor supply (Hall (1980)), in an econometric framework not nearly as fully worked out as the one used here. I am prepared to defend both conclusions on intuitive, practical grounds. People are quite willing to work hard this year and take it easy next year, in response to a modest incentive from real wages and real interest rates. They require much larger incentives to eat and drink more than usual this year and less than usual next year. Whatever cyclical fluctuations take place in consumption of nondurables (and they are very weak) probably cannot be attributed to the intertemporal substitution effects featured in modern theories of equilibrium business cycles. In fairness to the proponents of such theories, I don't think that intertemporal substitution in consumption has been given much of a role. The evidence of

this paper suggests we should stick with the labor supply side of household preferences in equilibrium explanations of fluctuations.

I plan a separate paper on some of the macroeconomic implications of low intertemporal substitutability of consumption, so I will confine myself to two brief comments. First, the substitution elasticity controls the speed of convergence of the simple general equilibrium model to its steady state. If the elasticity is zero, then convergence never occurs--the long-run state of the economy depends on its initial conditions. The simple idea conveyed by the model with positive substitution that eventually the economy moves to a point where the marginal product of capital equals the rate of time preference does not apply when the elasticity is zero. It is virtually irrelevant with very low but positive values of the elasticity, because convergence can take thousands of years.

Second, the strength of the intergenerational redistribution effects of the national debt or unfunded social security, debated recently by Robert Barro (1976) and Martin Feldstein (1976), depend on the elasticity of substitution. Of course, as Barro points out, if families behave as single individuals with infinite lifetimes, redistribution among generations is meaningless. But if the economy contains isolated individuals with finite lifetimes, then the elasticity of substitution governs the extent to which redistribution of consumption within lifetimes offsets the government's attempt to redistribute consumption across generations. With low substitution,

redistribution is highly effective--unfunded social security really does create more consumption for the older generation in general equilibrium.

References

Robert J. Barro, "Reply to Feldstein and Buchanan," Journal of Political Economy vol. 84, pp. 343-349, April 1976

Douglas T. Breeden, PhD dissertation, Graduate School of Business, Stanford University, December 1977

_____, "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," Journal of Financial Economics vol. 7, pp. 265-296, 1979

Eugene Fama, "Short-term Interest Rates as Predictors of Inflation," American Economic Review vol. 65, pp. 269-282, June 1975.

_____, "Stock Returns, Real Activity, Inflation and Money" University of Chicago School of Business, December 1980

Martin Feldstein, "Perceived Wealth in Bonds and Social Security," Journal of Political Economy vol. 84, pp. 331-336, April 1976

Sanford Grossman and Robert Shiller, "Consumption Correlatedness and Risk Measurement in Economies with Non-Traded Assets and Heterogeneous Information" University of Pennsylvania, Sloan Workshop in Transaction Cost Economics, April 1981

Robert E. Hall, "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," Journal of Political Economy, vol. 86, pp. 971-988, December 1978

_____, "Labor Supply and Aggregate Fluctuations," in K. Brunner and A. Meltzer (eds.) On the State of Macro-Economics, vol. 12 in the Carnegie-Rochester Conference Series on Public Policy, North-Holland, 1980

Lars Hansen and Kenneth Singleton, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Stock Market Returns," Carnegie-Mellon University, 1981

Robert E. Lucas, Jr., "Econometric Policy Evaluation: A Critique," in K. Brunner and A. Meltzer (eds.) The Phillips Curve and Labor Markets, volume 1 in the Carnegie-Rochester Conference Series on Public Policy, North-Holland, 1976

John Seater, "Marginal Federal Personal and Corporate Income Tax Rates in the U.S., 1909-1975," Research Papers of the Philadelphia Federal Reserve Bank, Number 57, November 1980