

THE SENSITIVITY OF CONSUMPTION TO TRANSITORY INCOME: ESTIMATES FROM PANEL DATA ON HOUSEHOLDS

BY ROBERT E. HALL AND FREDERIC S. MISHKIN¹

1. INTRODUCTION

THE STOCHASTIC RELATIONSHIP of consumption to income has long been recognized as a critical issue to macro policy analysis. One traditional view of consumers sees them as largely passive agents in the determination of aggregate demand. Changes in real incomes are translated reasonably quickly and fully into changes in consumption. In this view, income changes brought about by tax changes are a powerful tool for countercyclical stabilization policy, as Okun [14] and Tobin and Dolde [17] have argued. In contrast, the life cycle/permanent income hypothesis of consumption embodies the opposite view that consumers maximize utility over a long-term horizon. Rather than responding passively to every change in income, consumers will alter their consumption by smaller amounts if they perceive the income change as temporary rather than permanent. Eisner [3] has argued along this line. With the refinement of rational expectations, the life cycle/permanent income theory (as in Muth [13], Lucas [9], and Hall [7]) casts serious doubt on the usefulness of income variations as a stabilization tool. Consumers cannot be relied on to react vigorously when a policy-induced income change occurs. Predicting the impact of an income change on consumption requires knowledge of consumers' perceptions of its permanence.

This paper tries to shed some light on the stochastic relation between income and consumption (specifically, consumption of food) within a panel of about 2000 households who reported both variables over a seven-year span. Our major findings are:

1. Consumption responds much more strongly to permanent than to transitory movements of income.
2. The response to transitory income is nonetheless vigorous; it makes sense within the model only if interest rates are at least 20 per cent.
3. A simple test, independent of our model of consumption, rejects the pure life cycle/permanent income hypothesis.
4. The observed covariation of income and consumption is compatible with pure life cycle/permanent income behavior for 80 per cent of consumption and simple proportionality of consumption and income for the remaining 20 per cent.

These conclusions are derived from evidence about the joint movements of income and consumption. Needless to say, consumption and income frequently

¹This research was supported by the National Science Foundation through a grant to the National Bureau of Economic Research. We are grateful to numerous colleagues, especially Ariel Pakes, for helpful comments. The research reported here is part of the NBER's research program in Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

rise or fall together in the same year for a particular family. Our model explains the bulk of this correlation as the immediate response of consumption to changes in permanent income. Most of the rest is attributed to departures from life cycle/permanent income behavior for a small part of consumption. We hypothesize that about a fifth of all consumption is just set to a fraction of current income instead of following the more complicated optimal rule.

The nature of our data precludes any examination of intertemporal substitution effects in consumption. Thus our results do not answer the question of possible effects of temporary tax policies via substitution rather than income effects. Using time series data, Hall [8] finds that intertemporal substitution is apparently weak for total consumption of nondurable goods.

2. STOCHASTIC THEORY OF CONSUMER BEHAVIOR

An important paper by John Muth [13] on the permanent income hypothesis showed that the marginal propensities to consume out of current and lagged income depend on the stochastic properties of income. An income process with a large transitory component implies a small propensity to consume out of current income. At the other extreme, when most changes in income are permanent—that is, when income is almost a random walk—the propensity to consume out of current income should differ only slightly from the propensity to consume out of permanent income. This point was overlooked in empirical work on consumption long after the publication of Muth's article; Mayer's [11] survey does not mention any studies that consider the issue, for example. In recent work using data on individual consumers (Mayer [10]), estimates of large propensities to consume out of current income are interpreted as evidence against the permanent income hypothesis without any discussion of the stochastic process of income. The evidence is actually ambiguous because the permanent income hypothesis together with plausible income processes could well imply exactly the degree of sensitivity found.

Recent work by Hall [7] deals with some of these problems by deriving a theory of the stochastic process of consumption from the life cycle/permanent income hypothesis. Empirical tests then find that one of the important implications of the hypothesis is largely supported by aggregate time series data.² A recent paper by Flavin [5] examines aggregate data in a framework similar to the one used here, again with generally favorable results. However, aggregate evi-

²These tests and the empirical work here are closely related to the large body of research on the behavior of financial markets under rational expectations. Consumption is the analogue of a stock price, for example, and income is the analogue of earnings per share of the corporation issuing the stock. Our test of the predictive power of lagged income is the analogue of similar tests for market efficiency in the stock market, in the sense of the unpredictability of changes in stock prices from publicly available information (Fama [4] and Mishkin [12]). In contrast to our findings for consumption, the hypothesis of unpredictability is generally supported by the data for security markets. The technique developed in this paper could be transplanted directly to securities markets to answer such questions as: Do stock prices overreact to current movements of earnings? Do market participants have advance information about earnings?

dence is not really powerful enough to settle the important questions about the behavior of consumers.

These considerations have led to the research reported here based on data for individual households. We bring a rather specific question to this research: Are consumers more sensitive to current fluctuations in income than they would be if they followed the dictates of the life cycle/permanent income model? We approach the question in the following way: First, we propose a stochastic model of household income. Then, we hypothesize that households choose current consumption so as to maximize expected intertemporal utility, as suggested by the life cycle/permanent income view of consumption. In so doing, they arrive at an estimate of permanent income, based on the information available about the various stochastic components of actual income. Note that permanent income is not one of the components we hypothesize for actual income. Rather, permanent income is an intermediate step in the process by which families determine consumption. In this respect, we expand on earlier microeconomic research on the permanent income hypothesis. The final step makes observed consumption equal to a fraction of permanent income plus a transitory component which can be interpreted as measurement error, inventory accumulation, and the like.

The empirical analysis in this paper focuses on the theoretical implication that consumers should increase consumption by the annuity value of the increase in wealth brought about by a transitory increase in income. We test this implication by estimating the model using panel data on the income and consumption levels of individuals over several years. However, the response of consumption to the transitory component of income is estimated as a free parameter rather than constraining it to equal the expression for the annuity value. We then can evaluate whether consumption is excessively responsive to current income.

The starting point for our work is the life cycle/permanent income theory of consumption. According to the theory, consumers form estimates of lifetime resources and then adopt plans for spreading those resources over the remaining years of their lives. With explicit considerations of uncertainty (Yaari [18], Bewley [2], and Hall [7]), this principle becomes: Consumers form estimates of the probability distributions of lifetime resources and adopt sequential policies for spreading the resources. We will consider the hypothesis of rational expectations which asserts that consumers use all available information in estimating the probability distributions of future resources. This hypothesis is more of a sharpening and clarification of assumptions already implicit in the life cycle/permanent income theory rather than a logically independent assumption.

Here we consider the case of a household whose real income is the sum of three components: (i) a deterministic component, \bar{y}_t , which rises with age until just before retirement, and then falls rapidly; (ii) a stochastic component, y_t^L , which fluctuates as lifetime prospects change; because this lifetime component embodies information about essentially permanent family characteristics, a natural specification is a random walk; (iii) a stationary stochastic component y_t^S , which fluctuates according to transitory influences and is described by a moving average time-series process.

We use the following stochastic model of household income:

$$(1) \quad y_t = \bar{y}_t + y_t^L + y_t^S,$$

$$(2) \quad y_t^L = y_{t-1}^L + \epsilon_t,$$

$$(3) \quad y_t^S = \sum_{m=0}^{\infty} \phi_m \eta_{t-m} \quad (\phi_0 \text{ is normalized to equal one}),$$

where ϵ_t and η_t are random innovations that are completely unpredictable from past information. A key feature of our model is the hypothesis that families separately know the two stochastic components of income, ϵ_t and η_t .

With a quadratic utility function, $u(\tilde{c}_t) = -\frac{1}{2}(c_t^* - \tilde{c}_t)^2$ (where c_t^* is the bliss level of consumption and \tilde{c}_t is consumption), the household's intertemporal decision problem is to maximize

$$(4) \quad E_t \left[-\frac{1}{2} \sum_{\tau=0}^{T-t} (1+\delta)^{-\tau} (c_{t+\tau}^* - \tilde{c}_{t+\tau})^2 \right]$$

subject to the budget constraint³

$$(5) \quad \sum_{\tau=0}^{T-t} (1+r)^{-\tau} (y_{t+\tau} - \tilde{c}_{t+\tau}) + \tilde{A}_t = 0$$

where E_t is mathematical expectation conditional on all information available at time t ; δ is the rate of subjective time preference; r is real rate of interest which is assumed to be constant over time; T is length of economic life; \tilde{c}_t is consumption; c_t^* is bliss level of consumption; \tilde{A}_t is nonhuman wealth (assets).

The first order conditions for this problem are:

$$(6) \quad \frac{E_t[(1+\delta)^{-\tau} (c_{t+\tau}^* - \tilde{c}_{t+\tau})]}{c_t^* - \tilde{c}_t} = (1+r)^{-\tau} \quad \text{for } \tau = 1, \dots, T-t.$$

For a derivation, see Hall [7]. This equation states the intuitively plausible result that the marginal rate of substitution between current and future income equals the price ratio between current and future consumption.

In order to simplify the model used in estimation, we assume that the household's rate of time discount, δ , equals the real interest rate, r . Empirical evidence in Hall [7] does not call for rejection of this assumption, and it seems reasonable to make use of it in this paper's analysis. The first order conditions now imply

$$(7) \quad E_t \tilde{c}_{t+\tau} = c_{t+\tau}^* + \tilde{c}_t - c_t^*.$$

³Note that this budget constraint requires that y_t not include the expected return on nonhuman wealth, A_t . Flavin [5] discusses why the failure to recognize this can lead to errors.

Taking expectation conditional on information available at time t of the budget constraint and using (7) to substitute for $E_t \tilde{c}_{t+\tau}$ we have

$$(8) \quad \sum_{\tau=0}^{T-t} (1+r)^{-\tau} [E_t y_{t+\tau} - c_{t+\tau}^* - \tilde{c}_t + c_t^*] + \tilde{A}_t = 0.$$

Defining human wealth at time t as \tilde{H}_t ,

$$(9) \quad \tilde{H}_t \equiv \sum_{\tau=0}^{T-t} (1+r)^{-\tau} E_t y_{t+\tau},$$

and

$$(10) \quad \gamma_t \equiv \frac{1}{\sum_{\tau=0}^{T-t} (1+r)^{-\tau}},$$

$$(11) \quad \bar{V}_t \equiv \sum_{\tau=0}^{T-t} (1+r)^{-\tau} c_{t+\tau}^*,$$

we can rewrite (8) as

$$(12) \quad \tilde{c}_t = c_t^* + \gamma_t (\tilde{H}_t + \tilde{A}_t - \bar{V}_t).$$

This consumption function tells us that an increase of one unit of either human or nonhuman wealth will result in an increase of consumption of γ_t ; γ_t is easily recognizable as the annuity value of one unit of wealth, that is, the stream of equal payments over the remainder of the lifetime that can be financed by a unit of wealth. For infinite horizons and small r it approximately equals the real interest rate, r . Thus this consumption function is identical to the life cycle/permanent income consumption functions of Friedman [6] and Ando and Modigliani [1].

Further algebraic manipulation of the consumption function above is required to derive a model that can be estimated from the panel data used in this study. The deterministic paths of human and nonhuman wealth, denoted by \bar{H}_t and \bar{A}_t , respectively, are defined as the paths of these variables that occur in the absence of surprises (innovations) in y^L and y^S (specifically, $\bar{H}_t \equiv \sum_{\tau=0}^{T-t} (1+r)^{-\tau} \bar{y}_{t+\tau}$). The deterministic path of consumption, \bar{c}_t , is

$$(13) \quad \bar{c}_t = c_t^* + \gamma_t [\bar{H}_t + \bar{A}_t - \bar{V}_t].$$

We subtract (13) from (12) and define c_t and A_t as deviations from the deterministic paths of consumption and assets— $c_t = \tilde{c}_t - \bar{c}_t$ and $A_t = \tilde{A}_t - \bar{A}_t$ —to obtain

$$(14) \quad c_t = \gamma_t [H_t + A_t].$$

The evolution of assets around their deterministic path is governed by

$$(15) \quad A_t = (1 + r)(A_{t-1} + y_{t-1}^L + y_{t-1}^S - c_{t-1}).$$

Since the real return to assets, r , is assumed to be constant, A_t is completely known at $t - 1$. The deviation of human wealth around its deterministic path can be written as

$$\begin{aligned} (16) \quad H_t &= \sum_{\tau=0}^{T-t} (1 + r)^{-\tau} \left(y_t^L + \sum_{m=0}^{\infty} \phi_{\tau+m} \eta_{t-m} \right) \\ &= \frac{1}{\gamma_t} y_t^L + \sum_{m=0}^{\infty} \left(\sum_{\tau=0}^{T-t} (1 + r)^{-\tau} \phi_{\tau+m} \right) \eta_{t-m} \\ &= \frac{1}{\gamma_t} \epsilon_t + \frac{1}{\gamma_t} y_{t-1}^L + \left(\sum_{\tau=0}^{T-t} (1 + r)^{-\tau} \phi_{\tau} \right) \eta_t \\ &\quad + \sum_{m=1}^{\infty} \left(\sum_{\tau=0}^{T-t} (1 + r)^{-\tau} \phi_{\tau+m} \right) \eta_{t-m}. \end{aligned}$$

The model to be estimated could be derived through additional tedious algebra, but an easier route to this derivation makes use of a proposition demonstrated by Hall [7]. Consumers with rational expectations who maximize the expected value of an intertemporally separable utility function of the type outlined here will display the following condition: No information available to this consumer at $t - 1$ beyond the value of consumption will help in predicting next period's marginal utility of consumption. In the case of the quadratic utility function used here and the assumption that $r = \delta$, this implies that c_t is a random walk and Δc_t will equal the unpredictable component of c_t . Because A_t , η_{t-m} for $m \geq 1$, and y_{t-1}^L are all known at time $t - 1$, the only unpredictable elements of the wealth term $H_t + A_t$ in (14) are those which involve the contemporaneous innovations ϵ_t and η_t . Hence, the change in consumption is γ_t times the innovation in H_t :

$$(17) \quad \Delta c_t = \epsilon_t + \beta_t \eta_t$$

where

$$\beta_t = \gamma_t \sum_{\tau=0}^{T-t} (1 + r)^{-\tau} \phi_{\tau};$$

this is the annuity value of the added wealth implied by a unit amount of unexpected transitory income. Equation (17) is a convenient form that can be estimated with the panel data available to us because it eliminates the need for information on household nonhuman assets.

Our empirical analysis focuses on the coefficient β_t , which is the marginal propensity to consume out of the transitory increase in income, y_t . The derived value of β_t above shows that it should equal the annuity value of the addition to

wealth brought about by this transitory increase in income. Some illustrative values of β_t for the $MA(2)$ time series model of y_t^S estimated in this paper where the moving average parameters are .2 at lag one and .1 at lag two are as follows:

		For Real Interest Rate per Year Equal to			
		.05	.10	.20	.30
For Remaining	20 years	.095	.105	.170	.232
Lifetime	40 years	.071	.093	.167	.231

As these numbers indicate, rational consumption behavior is compatible with any degree of sensitivity to the surprise in the transitory income component (up to $\beta_t = 1$), provided a sufficiently high interest rate faces the consumer. No matter how much they discount the future, consumers should not simply make consumption proportional to current income; rather, the optimal strategy is to make the *change* in consumption respond only to the surprise in income and not to predictable movements of income. At very high interest rates, it is true that the information about future changes in income contained in today's surprise in income has negligible influence on wealth. However, it is still possible to take steps today that will insulate consumption from any foreseeable future changes in income. Exactly because the return to assets is high, a tiny amount saved from today's temporary increase in income can finance a complete offset of the subsequent decline in income later. In an economy with very high interest rates, consumers make small but lucrative and important asset transactions to achieve the optimal consumption path. Later in the paper, we will consider the behavior of consumers who are constrained against making any transactions in assets. They are prevented from achieving the optimal consumption path, and their actual consumption behaves in a way that is readily detectable in the data. There is a very substantial difference between optimal consumption in the face of very high interest rates and consumption constrained to equal current income.

As a final note on the interpretation of the theory, we emphasize that the lifetime component of income, y^L , is *not* the same thing as permanent income, although the propensity to consume out of y^L is the same as the propensity to consume out of permanent income. Permanent income includes the annuity values of transitory income and assets, as well as the lifetime component of income. Our research tries to make a clear distinction between the statistical decomposition of income into lifetime and transitory components, on the one hand, and the consumer's inference about permanent income, on the other hand.

3. STATISTICAL MODEL AND ESTIMATION

The data for our investigation are obtained from the University of Michigan's Panel Study of Income Dynamics (PSID) which contains histories of earnings

and spending for a large number of families over a span of several years. The PSID reports total annual family income net of estimated federal income taxes, which we then adjusted to take account of estimated FICA (social security) tax payments and changes in the overall cost of living (measured by the Consumer Price Index). The most comprehensive and reliable consumption measure which can be obtained from the PSID is the sum of the annual expenditures on food used at home and the amount spent eating at restaurants. We deflated food expenditures with the food price component of the CPI. Data from the PSID for food consumption are available for the years 1969–1971 and 1973–1975 and for income for all years, 1969–1975. We included all families who reported income and food consumption in all years and whose responses to the food and income questions were deemed accurate by the interviewer.⁴ We used data on six first differences of income and five first differences of consumption for 2309 families. One of the first differences of consumption spans two years; later in this section we describe how we accommodated this feature of the data.

In the PSID survey, information about food consumption is elicited by the question: “How much do you spend on food in an average week?” The question is asked sometime in the first half of the year; on the average the interview takes place at the end of March. We date the response in the previous year, as does the PSID. For a typical family interviewed in March 1971, for example, data on last year’s income and usual food consumption are dated 1970 in our work. Because of the peculiar timing of the question about average food consumption, we found it necessary to extend the model described earlier in the paper in the following way. We assume that the new information about income which the family uses to decide on consumption dated in year t includes a fraction ϕ of the new information that will not be recorded by the survey until the following year. For the simple reason that the consumption question is asked partway into the following year, we might expect a value of ϕ near a quarter. However, a family might have access to additional information about income for the full year at the time that consumption is measured early in the year. For example, in some jobs annual compensation is known with near certainty at the beginning of the year. If this kind of advance information about income is commonplace, our estimate of ϕ should be correspondingly higher. We do not consider the possibility that consumers have information about income in years after $t + 1$, beyond what can be predicted from the history of income itself. Our low estimated value of ϕ tends to confirm our assumption on this point. Further details about the role of future information in the model and the estimation of ϕ appear in the next section.

In addition to the ambiguity about the timing of the question about food consumption, there is further ambiguity about the length of the period over which consumption is measured. Instead of asking about average consumption over an unstated period, it would be better for our purposes if it were about last

⁴The survey interviewers were instructed to estimate income and food consumption when an interviewee was unsure of the answers to the questions concerning these items. We excluded all of the cases where this imputation was done.

year specifically or even about last week. We assume in the rest of this paper that the typical respondent averaged food consumption over a period much shorter than a year. In this case, the theory, which deals with consumption measured as an instantaneous flow, is a reasonable approximation for our data. Where consumption is measured as an average flow over an entire year, an explicit treatment of time aggregation is required.

The use of food consumption in place of total consumption obligates us to consider the form of the demand function for food, which differs in two respects from the demand function for total consumption. First, the price of food relative to the overall cost of living influences food consumption. Because all the families in the sample faced roughly the same change in relative prices, and our study relies primarily on the variability of individual family income, the relative price change presents few problems for our work. We posit equal relative price effects among families with similar characteristics, and remove these effects before estimating the model. Details of this adjustment appear later in the paper.

The second consideration is the likelihood that the proportion of income spent on food declines as income rises—the usual view about the Engel curve for food. In the current research, we approximate the Engel curve by a straight line with a positive intercept. Though this does imply a declining expenditure fraction on food, it can be defended only as an approximation. The slope of the line will be called α ; it is the marginal propensity to spend permanent income on food. The parameter β introduced in the previous section will be defined as the ratio of the marginal propensity to spend transitory income on food to the marginal propensity to spend lifetime income on food. Thus the units and the expected numerical values for β presented earlier will continue to apply.

Another extension of the basic model is necessary because food consumption is measured imperfectly. Any study of consumption at the level of individual households needs to include a stochastic element of measurement error and transitory consumption. We assume that measured consumption includes a transitory component, c_t^S , which obeys a second-order moving average process with parameters λ_1 and λ_2 :

$$(18) \quad c_t^S = v_t + \lambda_1 v_{t-1} + \lambda_2 v_{t-2}.$$

We hypothesize that transitory consumption is uncorrelated with both components of income:

$$(19) \quad \text{corr}(v_t, \epsilon_t) = \text{corr}(v_t, \eta_t) = 0.$$

With these various extensions, our model for the first difference of consumption becomes⁵

$$(20) \quad \Delta c_t = \alpha \epsilon_t + \alpha \beta \eta_t + v_t - (1 - \lambda_1) v_{t-1} - (\lambda_1 - \lambda_2) v_{t-2} - \lambda_2 v_{t-3}.$$

The terms involving v represent the first difference of a moving average process.

⁵ Here, we are neglecting the issue of advance information about income; the appropriate modifications are presented in the next section.

A detailed preliminary examination of the serial correlation properties of income revealed that a second-order moving average model was appropriate. With moving-average parameters ρ_1 and ρ_2 , the stochastic model for the first difference of income is

$$(21) \quad \Delta y_t = \epsilon_t + \eta_t - (1 - \rho_1)\eta_{t-1} - (\rho_1 - \rho_2)\eta_{t-2} - \rho_2\eta_{t-3}.$$

Again, the terms involving η are the first difference of a moving average process. This model embodies the strong assumption that income is measured without error. A model augmented with an income measurement error would not be econometrically identified.

Although in the full life cycle model, the propensity to consume out of transitory income depends on age, in the results presented here, we approximate the full model by treating β as constant across the sample. We tried estimating the model separately for families with younger and older heads, but failed to find significant differences. Constancy of β across families has the substantial statistical advantage of making the simple moment matrix over families a sufficient statistic for all of the parameters of the model.

We estimate the parameters of the model by maximum likelihood, under the assumption that ϵ , η , and ν obey normal distributions. Maximum likelihood achieves the best fit of the variances and covariances predicted by the model to those found in the data; the likelihood function is a scalar measure of the fit. The key idea of our approach is to write out the formulas for the variances and covariances of the data implied by our theoretical model, and then solve the resulting system of equations for the parameter estimates. To keep the exposition simple, we will first work out the case where transitory income and transitory consumption are not serially correlated (ρ_1 , ρ_2 , λ_1 , and λ_2 are all taken as zero) and no consumers have advance information about income ($\phi = 0$). First, the variance of the first difference of income is

$$(22) \quad V(\Delta y_t) = \sigma_\epsilon^2 + 2\sigma_\eta^2$$

and the covariance of the first difference of income with its own lagged value is

$$(23) \quad \text{cov}(\Delta y_t, \Delta y_{t-1}) = -\sigma_\eta^2.$$

These two formulas give us estimates of the variance of the innovation in transitory income, σ_η^2 , and of the variance of the increment in lifetime income, σ_ϵ^2 . Next, the covariance of the first difference of consumption with its own lagged value is

$$(24) \quad \text{cov}(\Delta c_t, \Delta c_{t-1}) = -\sigma_\nu^2.$$

This gives us the last of the three variances, that of the innovation in transitory consumption, σ_ν^2 .

Information about the structural parameters α and β comes from the covariances of consumption and income. The contemporaneous covariance is

$$(25) \quad C_0 = \text{cov}(\Delta c_t, \Delta y_t) = \alpha\sigma_\epsilon^2 + \alpha\beta\sigma_\eta^2$$

and the covariance with future income is

$$(26) \quad C_1 = \text{cov}(\Delta c_t, \Delta y_{t+1}) = -\alpha\beta\sigma_\eta^2.$$

Solving for α and β gives

$$(27) \quad \alpha = \frac{C_0 + C_1}{\sigma_\epsilon^2},$$

$$(28) \quad \beta = \frac{C_1}{C_0 + C_1} \frac{\sigma_\epsilon^2}{\sigma_\eta^2}.$$

It is not surprising that the contemporaneous covariance, C_0 , has a central role in estimating the two propensities to consume, α and β . It is perhaps a little surprising that the covariance of the current change in consumption with the future change in income is equally important. The basic finding of the paper is that this covariance is small, so it is not plausible that consumers are excessively sensitive to transitory income. Why would we expect excessive sensitivity to show up as a strong negative correlation between the change in consumption and the future change in income? Because those upward movements in consumption that are associated with the response to transitory income should be followed by a movement of income back toward normal in the following year. The first differences of income are negatively serially correlated (both in the theory and in the data), so the correlation of the change in consumption and the subsequent change in income should reflect this negative serial correlation.

It might appear that the covariance of current consumption and lagged income could provide similar information. That covariance is also free of the effects of changes in the lifetime component of income. However, the optimal use of information hypothesized for consumers in the model implies that the covariance should be exactly zero; no information available in year $t - 1$ should help predict the change in consumption in year t . This is essentially the proposition formulated and tested in Hall [7]. The test will be carried out with the micro panel data of this study in a later section of the paper.

4. ADVANCE INFORMATION ABOUT INCOME

Our data suggest that families have some information, but not full information, about next year's income innovations, ϵ_{t+1} and η_{t+1} , when they decide on this year's consumption, c_t . To incorporate this consideration, we introduce the future innovations into the consumption equation, weighted by a parameter, ϕ , which is interpreted as the fraction of advance information available to families. Because consumption is measured about a quarter of the way into the next year, a reasonable value for ϕ is around 0.25, and, indeed, our estimate is close to this value. The modified consumption equation is:

$$(29) \quad \Delta c_t = \alpha\phi\epsilon_{t+1} + \alpha(1 - \phi)\epsilon_t + \alpha\beta\phi\eta_{t+1} + \alpha\beta(1 - \phi)\eta_t + \Delta c_t^S.$$

This model, or one observationally equivalent, can be derived from either of two sets of assumptions:

ASSUMPTION 1: The annual income innovations ϵ_{t+1} and η_{t+1} are sums of, say, weekly innovations, and families have observed N of them when they make their consumption decisions; then $\phi = N/52$.

ASSUMPTION 2: Families observe a noisy value of the future innovation, and ϕ is the coefficient applied to the noisy value in forming the best forecast of the actual value.

For the simple case where the transitory components are serially uncorrelated ($\rho_1 = \rho_2 = \lambda_1 = \lambda_2 = 0$), the cross-covariances that identify ϕ and the other key parameters are

$$(30) \quad C_0 = \text{cov}(\Delta c_t, \Delta y_t) = (1 - \phi)(\alpha\sigma_\epsilon^2 + \alpha\beta\sigma_\eta^2),$$

$$(31) \quad C_1 = \text{cov}(\Delta c_t, \Delta y_{t+1}) = \phi(\alpha\sigma_\epsilon^2 + \alpha\beta\sigma_\eta^2) - (1 - \phi)\alpha\beta\sigma_\eta^2,$$

$$(32) \quad C_2 = \text{cov}(\Delta c_t, \Delta y_{t+2}) = -\phi\alpha\beta\sigma_\eta^2.$$

If ϕ is zero, the solution of these equations for α and β is the same as described in the previous section. If ϕ is one (complete advance information on income), the solution is the same, with C_1 taking the place of C_0 and C_2 taking the place of C_1 . In general, to solve for all three parameters, we start with

$$(33) \quad \alpha = \frac{C_0 + C_1 + C_2}{\sigma_\epsilon^2}.$$

The equations for β and ϕ are quadratic and it does not seem worth writing them out explicitly. Provided C_2 is negative (as it is in our data), the equations have a solution with ϕ between zero and one and a positive value of β .

In our model, we assume that families are homogeneous with respect to information about income—they all know a fraction ϕ of next year's income in making this year's consumption decisions. The covariances of this model are exactly the same as those for a model of heterogeneous families, where a fraction ϕ are fully aware of next year's income and the rest know nothing about it. The models are not completely the same, however. In the heterogeneous case, the distribution of Δc_t and Δy_t is not multivariate normal, but is a mixture of multivariate normals. Our estimates cannot be said to be maximum likelihood for the heterogeneous model.

The issue of advance information which might be available to market participants but not to the econometrician has also been considered in research on financial markets. The problem of the timing of the collection of data which obligates us to consider the issue here is not generally present in data on securities markets, but it may still be true that market participants have information in period t about what the econometrician labels an innovation in period

$t + 1$. One supporting piece of evidence is the predictive power of stock prices for future movements of the money stock, found by Rozeff [16] and Rogalski and Vinso [15].

5. ESTIMATION

Our bivariate model is

$$(34) \quad \Delta c_t = \alpha \phi \epsilon_{t+1} + \alpha(1 - \phi)\epsilon_t + \alpha\beta\phi\eta_{t+1} \\ + \alpha\beta(1 - \phi)\eta_t + \nu_t - (1 - \lambda_1)\nu_{t-1} - (\lambda_1 - \lambda_2)\nu_{t-2} - \lambda_2\nu_{t-3}$$

and

$$(35) \quad \Delta y_t = \epsilon_t + \eta_t - (1 - \rho_1)\eta_{t-1} - (\rho_1 - \rho_2)\eta_{t-2} - \rho_2\eta_{t-3}.$$

Let x be the column vector of unobserved random variables,

$$(36) \quad x' = [\epsilon_1, \dots, \epsilon_7, \nu_{-2}, \dots, \nu_7, \eta_{-2}, \dots, \eta_7].$$

We assume that x is multivariate normal, with a diagonal covariance matrix, Σ , and variances

$$(37) \quad V(\epsilon_t) = \sigma_\epsilon^2,$$

$$(38) \quad V(\nu_t) = \sigma_\nu^2,$$

$$(39) \quad V(\eta_t) = \sigma_\eta^2.$$

Let z_t be the column vector containing the 5 differences of consumption and 6 first differences of income for family i :

$$(40) \quad z_i' = [\Delta c_1, \Delta c_2, \Delta c_3 + \Delta c_4, \Delta c_5, \Delta c_6, \Delta y_1, \Delta y_2, \Delta y_3, \Delta y_4, \Delta y_5, \Delta y_6].$$

The model can be stated in the form

$$(41) \quad z_i = Ax_i$$

(note that the third row of A has a special form) and z_i is multivariate normal. The covariance matrix of z_i is

$$(42) \quad \Omega(\theta) = A\Sigma A'.$$

Here θ is the vector of parameters,

$$(43) \quad \theta' = [\alpha, \beta, \lambda_1, \lambda_2, \rho_1, \rho_2, \sigma_\epsilon^2, \sigma_\nu^2, \sigma_\eta^2, \phi].$$

The log-likelihood of the sample is

$$(44) \quad L(\theta) = -\frac{N}{2} \log \det \Omega(\theta) - \frac{1}{2} \sum_{i=1}^N z_i' \Omega^{-1}(\theta) z_i$$

plus an inessential constant. We estimate θ by full numerical maximization of the likelihood. Its estimated covariance matrix is computed as the inverse of the information matrix, $\partial^2 L / \partial \theta \partial \theta'$. All computations were carried out by a program written by Bronwyn Hall, which uses analytical derivatives and the method of scoring.

6. DETERMINISTIC COMPONENTS OF INCOME AND FOOD CONSUMPTION

The data whose variances and covariances are the starting point for the estimation process are the deviations of the changes in food consumption and income from deterministic paths. To form the deviations, we need estimates of the deterministic changes in income and consumption for each family in each year based on the family's characteristics in that year. We do this by assuming that the deterministic changes are functions of the family characteristics, then use ordinary least-squares regressions as follows: In the case of income, we regress the change in actual income on an intercept, the age of the household head, the age of the household head squared, the change in the number of adults in the household, the change in the number of children in the household, and a linear time trend. Since food is a commodity whose relative price changed substantially over the period of our sample, we need to take account of the downward slope of families' demand functions. Thus the change in food consumption is regressed on the percentage change in the relative price of food (as measured by CPI components) as well as on the variables used in the income regressions. Results for the income and food consumption regressions are

$$\begin{aligned} \Delta \text{INCOME} = & - \frac{433.96}{(182.08)} + \frac{33.23}{(7.89)} \text{AGE} - \frac{.35}{(.082)} \text{AGE}^2 \\ & + \frac{504.07}{(30.55)} \Delta \text{CHILD} + \frac{1535.06}{(40.99)} \Delta \text{ADULT} \\ & - \frac{53.44}{(13.04)} \text{TIME}, \end{aligned}$$

13854 observations, $R^2 = .1383$, Standard Error = \$2606.4;

$$\begin{aligned} \Delta \text{FOOD} = & - \frac{96.67}{(32.38)} + \frac{3.89}{(1.32)} \text{AGE} - \frac{.045}{(.014)} \text{AGE}^2 \\ & + \frac{166.56}{(6.39)} \Delta \text{CHILD} + \frac{242.46}{(8.72)} \Delta \text{ADULT} \\ & + \frac{2.00}{(2.65)} \text{TIME} - \frac{440.62}{(244.85)} \Delta \text{LOG PRICE}, \end{aligned}$$

11545 observations, $R^2 = .1438$, Standard Error = \$542.02;

where ΔINCOME is the change in family income which is adjusted for income and FICA taxes and the cost of living, ΔFOOD is the change in family spending for food at home and in restaurants deflated into real terms, AGE is age of household head, AGE^2 is AGE squared, ΔCHILD is the change in the number of children in the household, ΔADULT is the change in the number of adults in the household, TIME is time trend ($1970 = 1 \cdots 1975 = 6$), $\Delta \text{LOG PRICE}$ is the change in the log of the relative price of food (measured by the food component of the CPI deflated by the overall CPI), and standard errors of the coefficients are in parentheses. The residuals from these regressions are then taken to be the deviations from the deterministic paths of changes in food consumption and income.

The specification of the food and income regressions make little difference to the results obtained for the stochastic model outlined above. For example, if the effect of family characteristics on the deterministic paths of income and food consumption are ignored—i.e., the change in the deterministic components is just assumed to be a constant—we find only very small differences in the estimates of the parameters of the stochastic model. For this reason, we believe that our implicit assumption that families have perfect foresight about the variables on the right-hand side of these regressions is an innocuous one, though clearly overly strong.

7. RESULTS

The residuals from the preliminary regressions showed mild heteroskedasticity, especially in the first difference of consumption. Rather than complicate the model by introducing separate variances for each year, we simply transformed the covariance matrix of the residuals by dividing its rows and columns by suitable constants so that the variances of the first differences of consumption were the same in all years (equal to the average of the original data over the same years). We applied the same transformation to the income data. The spirit of this preliminary treatment of the data is the same as conversion to a correlation matrix, but it preserves the units of the structural parameters. Experiments with the alternative of estimating variances gave essentially the same estimates of the structural parameters.

Estimation by maximum likelihood yielded the results shown in Table I. In summary, they show:

1. The marginal propensity to consume lifetime income on food, α , is about 0.11, well under the average propensity in the raw data of 0.19.
2. The marginal propensity to consume out of transitory income relative to the marginal propensity to consume out of lifetime income, β , is estimated as 0.29, somewhat above its theoretical value at reasonable discount rates. The hypothesis of equal response to both components, $\beta = 1$, is unambiguously rejected.
3. The fraction of information about next year's income, ϕ , is 0.25, in line with prior expectations.

TABLE I
RESULTS FOR BASIC MODEL

Parameter	Value (Standard Error)	Interpretation
α	.107 (.008)	Fraction of permanent income spent on food
β	.292 (.080)	Relative effect of innovation in transitory income compared to effect of innovation in lifetime income
ϕ	.253 (.058)	Fraction of information available in year t about income in year $t + 1$
λ_1	.215 (.014)	First moving average parameter for transitory consumption
λ_2	.101 (.017)	Second moving average parameter for transitory consumption
ρ_1	.294 (.021)	First moving average parameter for transitory income
ρ_2	.114 (.018)	Second moving average parameter for transitory income
σ_ϵ^2	1.49 (.11)	Variance of innovation in lifetime income (thousands of dollars squared)
σ_r^2	.158 (.003)	Variance of innovation in transitory consumption
σ_η^2	3.41 (.13)	Variance of innovation in transitory income

Table II presents a reasonably complete accounting of the success of the model in fitting the pattern of covariation found in the data. For estimation of the key parameters α , β , and ϕ , the covariances of this year's change in consumption with this year's change in income, next year's change, and the subsequent year's change are the most important. All three parameters control the fitted value of the contemporaneous covariance— α and β make it larger, by making the change in consumption more sensitive to surprises in income, while ϕ makes it smaller, by making part of this year's change in consumption depend on next year's surprise in income. For the covariance with next year's income, β makes the fitted value more negative, for the reason explained earlier—if this year's consumption is sensitive to this year's transitory income, it will be negatively related to the change in next year's income when the transitory movement will probably be reversed. On the other hand, the fitted covariance is positively related to ϕ . If consumption is partly based on information about next year's surprise in income, this year's change will be positively correlated with

TABLE II
ACTUAL AND FITTED COVARIANCES

	Actual	Fitted
$\text{Var}(\Delta c)$.285	.285
$\text{Var}(\Delta y)$	6.772	6.757
$\text{Cov}(\Delta c, \Delta y)$.234	.200
$\text{Cov}(\Delta c, \Delta y_{+1})$	-.004	.003
$\text{Cov}(\Delta c, \Delta y_{+2})$	-.021	-.038
$\text{Cov}(\Delta c, \Delta y_{-1})$	-.077	.000
$\text{Cov}(\Delta c, \Delta c_{-1})$	-.110	-.106
$\text{Cov}(\Delta y, \Delta y_{-1})$	-1.948	-1.904
$\text{Cov}(\Delta y, \Delta y_{-2})$	-.319	-.339
$\text{Cov}(\Delta y, \Delta y_{-3})$	-.383	-.389

NOTES: $\text{Var}(\Delta c)$ includes $\text{var}(\Delta c_3 + \Delta c_4)$; $\text{cov}(\Delta c, \Delta y)$ includes $\text{cov}(\Delta c_3 + \Delta c_4, \Delta y_3)$ and $\text{cov}(\Delta c_3 + \Delta c_4, \Delta y_4)$; $\text{cov}(\Delta c, \Delta y_{+1})$ includes $\text{cov}(\Delta c_3 + \Delta c_4, \Delta y_5)$; $\text{cov}(\Delta c, \Delta y_{+2})$ includes $\text{cov}(\Delta c_3 + \Delta c_4, \Delta y_6)$; $\text{cov}(\Delta c, \Delta y_{-1})$ includes $\text{cov}(\Delta c_3 + \Delta c_4, \Delta y_2)$, and $\text{cov}(\Delta c, \Delta c_{-1})$ includes $\text{cov}(\Delta c_3 + \Delta c_4, \Delta c_2)$ and $\text{cov}(\Delta c_5, \Delta c_3 + \Delta c_4)$.

next year's change in income. The fitted covariance of almost exactly zero represents cancellation of the two effects, since both β and ϕ are quite positive. The estimation process separates the effects of β and ϕ through the use of the covariance of this year's change in consumption with the change in income two, three, four, and five years from now (of these, the closer ones are relatively more important). Under the hypothesis implicit in our model that consumers have no information about surprises in income more than one year in advance, the only explanation of the negative covariation of current consumption and future income operates through the sensitivity of consumption to transitory income, controlled by β . The estimation process chooses a substantially positive value of β in order to try to match the covariance of $-.021$; the overstatement in the fitted value of $-.038$ corresponds to understatements of some of the more distant covariances not shown in Table II.

An alternate explanation of the negative correlation of Δc_t and Δy_{t+2} is the possible inability of families to distinguish innovations in lifetime income from innovations in transitory income. If they cannot make the distinction at all, they are forced to react equally to both innovations, and so the estimate of β would be close to one. Our finding of a β above the level suggested by the theory may be a sign of limited information, not a sign of irrational behavior.

The only serious failure of the model revealed in Table II is its inability to explain the observed negative correlation of the current change in consumption and the lagged change in income. As we will show, the actual correlation is statistically significantly negative, yet the model holds that it should be exactly zero. The theoretical justification for the fitted correlation of zero is simple: Apart from its transitory component, consumption should respond only to new information, and lagged income cannot contain any new information. The next section of the paper examines the apparent failure of this principle.

8. THE RELATION BETWEEN CONSUMPTION AND LAGGED INCOME

The model has the straightforward implication that the simple regression of Δc_t on Δy_{t-1} should yield a coefficient on Δy_{t-1} of zero. Instead we find

$$\Delta c_t = - \frac{4.95}{(6.16)} - \frac{0.010}{(.002)} \Delta y_{t-1},$$

$$6926 \text{ observations,} \quad \text{standard error} = \$512, \quad R^2 = .0028.$$

Though the coefficient is quite small, it is statistically unambiguously negative. It would be uninteresting to conclude that the measurement error in Δc_t was negatively correlated with Δy_{t-1} , so we restrict our attention to explanations of a negative relation between the true change in consumption and the lagged change in income.

In this section we investigate the possibility that consumers are actually more sensitive to transitory income than is predicted by theory, but in a way not revealed in our examination of the joint behavior of Δc_t and Δy_t . The results in the previous section did not draw on the observed correlation between Δc_t and Δy_{t-1} —maximum likelihood is blind to covariances whose theoretical values are zero for all values of the parameters. An extended model proposed in Hall [7] for a similar purpose can be used to examine the lagged relation. Suppose that a fraction $1 - \mu$ of consumption follows the life cycle/permanent income theory and the rest, a fraction μ , simply tracks current total income passively and so has an excessive sensitivity to transitory income. If all consumption were simply proportional to current income, the model would take the form,

$$\begin{aligned} (45) \quad \Delta c_t &= \alpha \Delta y_t + \Delta c_t^S \\ &= \alpha \epsilon_t + \alpha \eta_t - \alpha(1 - \rho_1)\eta_{t-1} - \alpha(\rho_1 - \rho_2)\eta_{t-2} - \alpha\rho_2\eta_{t-3} \\ &\quad + v_t - (1 - \lambda_1)v_{t-1} - (\lambda_1 - \lambda_2)v_{t-2} - \lambda_2v_{t-3}. \end{aligned}$$

The covariance of the change in consumption with last year's change in income implied by this model is

$$(46) \quad \text{cov}(\Delta c_t, \Delta y_{t-1}) = -\alpha[(1 - \rho_1 + \rho_2)^2 - \rho_2]\sigma_\eta^2$$

which is negative. The logic of the negative covariance is straightforward: If consumption tracks income, then a transitory rise in income this year will typically be followed next year by a decline in income and so also in consumption.

We estimated a model in which a fraction μ of consumption tracks income (measured as ϕ times next year's income plus $1 - \phi$ times this year's income) and a fraction $1 - \mu$ responds in the way suggested by the life cycle/permanent

income hypothesis. The model is

$$(47) \quad \Delta c_t = \alpha(\phi\epsilon_{t+1} + (1 - \phi)\epsilon_t) + \mu\alpha(\phi\Delta y_{t+1}^S + (1 - \phi)\Delta y_t^S) \\ + (1 - \mu)\alpha\beta(\phi\eta_{t+1} + (1 - \phi)\eta_t) + \Delta c_t^S,$$

$$(48) \quad \Delta y_t = \epsilon_t + \Delta y_t^S \\ = \epsilon_t + \eta_t - (1 - \rho_1)\eta_{t-1} - (\rho_1 - \rho_2)\eta_{t-2} - \rho_2\eta_{t-3}.$$

The results of estimating the augmented model are:

α	.100 (.009)	propensity to consume food out of lifetime income;
β	.174 (.100)	propensity to consume out of a transitory increase in income relative to propensity to consume out of lifetime income;
ϕ	.206 (.058)	fraction of information available in year t about income in year $t + 1$;
μ	.200 (.065)	fraction of consumption directly proportional to current income.

The other parameter estimates are similar to their previous values. The new specification is about halfway successful in matching the covariance of this year's change in consumption with last year's change in income—the predicted value is $-.032$ against the sample value of $-.077$. Not surprisingly, the sensitivity of consumption to the innovation in transitory income is found to be smaller in the extended model, as the positive estimate of μ has taken over part of the job of explaining the positive contemporaneous covariation of consumption and income. Further, because μ and β are partly estimated from the same features of the data, joint estimation very substantially raises the sampling variation of the estimate of β , relative to the earlier results. The confidence interval for β now includes the theoretically expected value of about 0.10.

9. CONCLUDING REMARKS

According to our extended model, about 80 per cent of consumption obeys the life cycle/permanent income hypothesis. Consumption does not adjust in the same mechanical way to every change in income. Instead, consumers think about the source of a change in income and react vigorously only to those changes that signal a major shift in economic well-being. But the data reject the strong hypothesis that all consumption is governed by the life cycle/permanent income

principle. This conclusion is independent of the model developed in this paper; it rests solely on the rather general principle that changes in consumption should not be predictable on the basis of information available to the household. The negative relation between the lagged change in income and the current change in consumption is consistent with constrained consumption behavior for about 20 percent of consumption. We are able to distinguish this symptom of inability (or unwillingness) to borrow and lend from the type of behavior characteristic of consumers who simply face high effective interest rates. The data show signs of both influences. Consumption is somewhat more sensitive to current income than it would be in an economy where every consumer borrowed and lent freely at the Treasury bill rate. Still, it is much less sensitive than in an economy where no consumer ever borrowed or lent at all.

The overwhelming bulk of the movements in income that give rise to our inference from the data are unrelated to the behavior of the national economy; most are probably highly personal. It is purely an inference, though a reasonable one in our opinion, that households respond to income fluctuations attributable to the business cycle or to countercyclical tax policy in the same way they respond to purely personal income fluctuations. Our results cast doubt on the wisdom of tax policies to manipulate aggregate demand by changing disposable income. If, as the results indicate, most consumers react only to the new information about their permanent incomes conveyed by the announcement of a tax change, then policy-makers face the complicated task of inferring consumers' interpretation of the announcement. Lucas [9] has pointed out the obstacles to policy evaluation in these circumstances.

Our evidence and conclusions refer specifically to food consumption and more generally to the consumption of nondurables. Nothing in our work describes the response of consumer purchases of durable goods to changes in income. Our finding that food consumption behaves as if constraints on borrowing were relatively unimportant does not rule out important constraints for the acquisition of durables. The sensitivity of durables purchases to transitory income is very definitely a topic for further research, where some of the techniques developed in this paper may be helpful.

Stanford University
and
University of Chicago

Manuscript received April, 1980; revision received March, 1981.

REFERENCES

- [1] ANDO, A., AND F. MODIGLIANI: "The 'Life-Cycle' Hypothesis of Saving: Aggregate Implications and Tests," *American Economic Review*, 53(1963), 55-84.
- [2] BEWLEY, T.: "The Permanent Income Hypothesis: A Theoretical Formulation," Department of Economics, Harvard University, September, 1976.
- [3] EISNER, R.: "Fiscal and Monetary Policy Reconsidered," *American Economic Review*, 59(1969), 897-905.

- [4] FAMA, E. F.: "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance*, 25(1970), 383–417.
- [5] FLAVIN, M.: "The Adjustment of Consumption to Changing Expectations About Future Income," *Journal of Political Economy*, 89(1981), 974–1009.
- [6] FRIEDMAN, M.: *A Theory of the Consumption Function*. Princeton, New Jersey: Princeton University Press, 1957.
- [7] HALL, R. E.: "Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86(1978), 971–987.
- [8] ———: "Consumption and Real Interest Rates," (processed, 1980).
- [9] LUCAS, R. E.: "Econometric Policy Evaluation: A Critique," in *The Phillips Curve and Labor Markets*, Carnegie-Rochester Conference Series on Public Policy, 1, ed. by K. Brunner and A. H. Meltzer. Amsterdam: North Holland, 1976, pp. 19–46.
- [10] MAYER, T.: "Tests of the Permanent Income Theory with Continuous Budgets," *Journal of Money Credit and Banking*, 6(1972), 757–778.
- [11] ———: *Permanent Income, Wealth and Consumption*. Berkeley: University of California Press, 1972.
- [12] MISHKIN, F. S.: "Efficient-Markets Theory: Implications for Monetary Policy," *Brookings Papers on Economic Activity*, 2(1978), 707–752.
- [13] MUTH, J. F.: "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, 55(1960), 299–306.
- [14] OKUN, A.: "The Personal Tax Surcharge and Consumer Demand," *Brookings Papers on Economic Activity*, 1(1971), 167–204.
- [15] ROGALSKI, R. J., AND J. D. VINSO: "Stock Returns, Money Supply and the Direction of Causality," *Journal of Finance*, 32(1977), 1017–1030.
- [16] ROZEFF, M. S.: "Money and Stock Prices: Market Efficiency and the Lag in the Effect of Monetary Policy," *Journal of Financial Economics*, 1(1974), 245–302.
- [17] TOBIN, J., AND W. DOLDE: "Wealth, Liquidity and Consumption," in *Consumer Spending and Monetary Policy: The Linkages*, Federal Reserve Bank of Boston, 1971, 99–146.
- [18] YAARI, M.: "A Law of Large Numbers in the Theory of Consumer's Choice under Uncertainty," *Journal of Economic Theory*, 12(1976), 202–217.

