

ALAN S. BLINDER

Brookings Institution and Princeton University

ANGUS DEATON

Princeton University

The Time Series Consumption Function Revisited

THE RELATIONSHIP between consumer spending and income is one of the oldest statistical regularities of macroeconomics—and one of the sturdiest. Like the aging movie star, it needs a little touching up now and again, but always seems to come bouncing back.

A dozen years ago, both the theoretical derivation and the econometric form of the aggregate consumption function were considered settled. Most economists adhered to one of two ways of putting Fisher's theory of intertemporal optimization into operation: Milton Friedman's permanent income hypothesis (henceforth, PIH) or Franco Modigliani's life-cycle hypothesis (henceforth, LCH).¹ Since each variant seemed to have sound theoretical underpinnings, and since the two had similar econometric forms that explained the data well and had similar implications for policy, there was not a great deal to quarrel about. Perhaps the most contentious empirical issue was the apparently large marginal

This paper has benefited from the comments and suggestions of Albert Ando, Whitney Newey, and members of the Brookings Panel and from seminar presentations at the University of Warwick, Princeton University, and Johns Hopkins University. We thank Peter Rathjens and Lori Grunin for research assistance and the National Science Foundation for financial support.

1. Milton Friedman, *A Theory of the Consumption Function* (Princeton University Press, 1957); Franco Modigliani and Richard Brumberg, "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in Kenneth K. Kurihara, ed., *Post-Keynesian Economics* (Rutgers University Press, 1954), pp. 388–436; Albert Ando and Franco Modigliani, "The 'Life-Cycle' Hypothesis of Saving: Aggregate Implications and Tests," *American Economic Review*, vol. 53 (May 1963), pp. 55–84.

propensity to consume out of transitory income, which was variously explained by a “short horizon” (that is, a high discount rate) or by liquidity constraints.

Things are quite different now. Developments in economic research, as well as actual events, have raised fundamental questions about the consumption function. At the same time, the range of experience of the last dozen years has been great enough to hold out the hope of getting some answers from aggregate data. This seems, therefore, an auspicious time to take a fresh, and unabashedly empirical, look at the time series consumption function.

Questions Raised by Modern Research

The Lucas Critique. Reasons abound for questioning the traditional consumption function and its implications for how tax policy affects consumer spending. Robert Lucas has pointed out that, under rational expectations, the PIH does not lead to a “structural” relationship between consumption and income, but rather to a statistical relationship that should change whenever the stochastic process generating income changes. The Lucas critique calls for estimation methods that treat consumption and income jointly.²

The “Random Walk” Hypothesis. Robert Hall sharpened the implications of the PIH by showing that the rational expectations hypothesis implies that only “surprises” in permanent income should affect current consumption, once lagged consumption is controlled for.³

The work of Hall and Lucas added a new dichotomy—that between anticipated and unanticipated changes in income—to the traditional permanent-transitory dichotomy. It is this new dichotomy, rather than the old one, that has absorbed the attention of contemporary researchers. Hall’s work in particular has spawned an infant industry estimating Euler equations linking current and lagged consumption in the manner

2. Robert E. Lucas, Jr., “Econometric Policy Evaluation: A Critique,” in Karl Brunner and Allan H. Meltzer, eds., *The Phillips Curve and Labor Markets*, Carnegie-Rochester Conference Series on Public Policy, vol. 1 (Amsterdam: North-Holland, 1976), pp. 19–46.

3. Robert E. Hall, “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy*, vol. 86 (December 1978), pp. 971–87.

implied by the first-order conditions of a Fisherian intertemporal optimization problem. We have our doubts about the wisdom of modeling aggregate consumption as the interior solution to a single individual's optimization problem in adjacent periods,⁴ but in any case think it fair to say that the research done to date has not supported the econometric restrictions implied by the Euler equation approach. Nor has further investigation validated the hypothesis that the response of consumption to income (henceforth, Y) reflects only the usefulness of current Y in predicting future Y . Instead, research typically finds "excess sensitivity" to current income.⁵ But the case is by no means closed. So a central question of this study is whether information known at time $t-1$, such as anticipated income, has any predictive power for changes in consumption between times $t-1$ and t .

The Barro Equivalence Hypothesis. A rather different objection to standard consumption functions, based on the idea that private and government accounts should be consolidated, was raised by Robert Barro.⁶ The income (that is, disposable income) and wealth (that is, household net worth, including government debt) variables normally used in consumption functions imply that intertemporal shifts in the

4. Corner solutions stemming from liquidity constraints pose one problem. Aggregation poses others. For example, with mortal consumers, a constant age distribution of the population, and neither surprises nor changes in interest rates, the ratio C_{t+1}/C_t (where C denotes consumption) would be 1 plus the growth rate of per capita income. The Euler equation approach models the growth rate of consumption as a function of the interest rate and the time discount rate, and uses the observed growth rate to estimate the intertemporal elasticity of substitution. Background growth of per capita income seems to be ignored.

5. Marjorie Flavin, "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy*, vol. 89 (October 1981), pp. 974-1009; Robert E. Hall and Frederic S. Mishkin, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, vol. 50 (March 1982), pp. 461-81; Ben S. Bernanke, "Adjustment Costs, Durables, and Aggregate Consumption," *Journal of Monetary Economics*, vol. 15 (January 1985), pp. 41-68; Ben S. Bernanke, "Permanent Income, Liquidity, and Expenditure on Automobiles: Evidence from Panel Data," *Quarterly Journal of Economics*, vol. 99 (August 1984), pp. 587-614; N. Gregory Mankiw, Julio J. Rotemberg, and Lawrence H. Summers, "Intertemporal Substitution in Macroeconomics," *Quarterly Journal of Economics*, vol. 100 (February 1985), pp. 225-51; Martin Browning, Angus Deaton, and Margaret Irish, "A Profitable Approach to Labor Supply and Commodity Demands over the Life Cycle," *Econometrica*, vol. 53 (May 1985), pp. 503-43; Charles R. Bean, "The Estimation of 'Surprise' Models and the 'Surprise' Consumption Function," Discussion Paper 54 (Center for Economic Policy Research, February 1985).

6. Robert J. Barro, "Are Government Bonds Net Wealth?" *Journal of Political Economy*, vol. 82 (November-December 1974), pp. 1095-1117.

pattern of taxes, with no change in their present value, produce shifts in the time pattern of consumption. This should not be so, Barro argued, if people can freely transfer income across generations.

The Barro equivalence hypothesis is not theoretically unobjectionable. In addition to the usual perfect capital markets assumption, it requires that bequests be motivated by intergenerational altruism and that people have extremely long time horizons. It also has trouble dealing with childless people or with the possibility of "corner solutions" in which the unconstrained optimal bequest cannot be enforced because it is negative. Because of these and other problems, many economists find the equivalence hypothesis implausible on a priori grounds. But a priori reasoning is not the way to settle the issue, and empirical studies have found it surprisingly difficult to reject the equivalence hypothesis.⁷ More evidence would be welcome, and we try to obtain some below.

Intertemporal Substitution. Modern macroeconomic analysis has reemphasized intertemporal substitution. Yet standard consumption functions often omit the rate of interest as an argument—not on theoretical grounds, but on empirical grounds. The consensus conclusion that consumption, and hence saving, is insensitive to the rate of return has been questioned by Michael Boskin and, more recently, by Lawrence Summers.⁸ What do recent data say about this issue?

7. For some theoretical arguments, see Martin S. Feldstein, "Perceived Wealth in Bonds and Social Security: A Comment," *Journal of Political Economy*, vol. 84 (April 1976), pp. 331–36; Robert J. Barro, "Reply to Feldstein and Buchanan," *Journal of Political Economy*, vol. 84 (April 1976), pp. 343–49; and Willem H. Buiter and James Tobin, "Debt Neutrality: A Brief Review of Doctrine and Evidence," in George M. von Furstenberg, ed., *Social Security vs. Private Saving* (Ballinger, 1979), pp. 39–64. The conclusion that the empirical evidence is mixed is reached by Karl Brunner after a thorough review of the literature. See Brunner, "Fiscal Policy in Macro Theory: A Survey and Evaluation" (Federal Reserve Bank of St. Louis, January 1985). But both Roger C. Kormendi, in "Government Debt, Government Spending, and Private Sector Behavior," *American Economic Review*, vol. 73 (December 1983), pp. 994–1010, and John J. Seater and Roberto S. Mariano, in "New Tests of the Life Cycle and Tax Discounting Hypotheses," *Journal of Monetary Economics*, vol. 15 (March 1985), pp. 195–215, claim that the data strongly support the Barro hypothesis. For a recent contrary view, see Michael J. Boskin and Laurence J. Kotlikoff, "Public Debt and U.S. Saving: A New Test of the Neutrality Hypothesis," Working Paper 1646 (National Bureau of Economic Research, June 1985).

8. Michael J. Boskin, "Taxation, Saving, and the Rate of Interest," *Journal of Political Economy*, vol. 86 (April 1978), part 2, pp. S3–S27, and Lawrence H. Summers, "Tax Policy, the Rate of Return, and Savings" Working Paper 995 (National Bureau of Economic Research, September 1982). For a critique of Boskin's work and opposing results,

Questions Raised by Recent Events

Temporary Tax Changes. The pure PIH with no liquidity constraints predicts that people will react much less to temporary than to permanent changes in taxes. And in 1968, when a temporary tax surcharge was imposed, consumption did indeed decline less than simple Keynesian consumption functions predicted. Similarly, a temporary tax decrease in 1975 led to a strong surge in saving, but only a modest increase in consumption. In the aftermath of these two episodes, both casual observation of the facts and formal econometric research seemed to support a modified version of the PIH.⁹ But this inference rested on a slender data base.

Recent events have given us another episode. Since the Reagan tax cuts of 1981–84 came in three preannounced stages, they can be thought of as a permanent tax reduction in August 1981 coupled with a temporary tax increase of gradually diminishing size. Thus the PIH predicts that saving should have declined sharply after August 1981 as the scheduled permanent tax cut induced higher consumption. Did it?

To answer this question we must, at a minimum, adjust the data for the sharp business cycle that took place during this period, for even very weak versions of the PIH imply that saving rates should fall in downturns and rise in booms. Table 1 shows cyclically adjusted saving rates for

see E. Philip Howrey and Saul H. Hymans, "The Measurement and Determination of Loanable-Funds Saving," *BPEA*, 3:1978, pp. 655–85. Robert Hall, in "Intertemporal Substitution in Consumption" (Stanford University, July 1985), argues that the interest elasticity of consumption is quite small, and that previous high estimates are biased. For two recent empirical surveys, see Thorvaldur Gylfason, "Interest Rates, Inflation, and the Aggregate Consumption Function," *Review of Economics and Statistics*, vol. 63 (May 1981), pp. 233–45; and Gerald Carlino, "Interest Rate Effects and Intertemporal Consumption," *Journal of Monetary Economics*, vol. 9 (March 1982), pp. 223–34.

9. Robert Eisner, "Fiscal and Monetary Policy Reconsidered," *American Economic Review*, vol. 59 (December 1969), pp. 897–905; Arthur M. Okun, "The Personal Tax Surcharge and Consumer Demand, 1968–1970," *BPEA*, 1:1971, pp. 167–204; William L. Springer, "Did the 1968 Surcharge Really Work?" *American Economic Review*, vol. 65 (September 1975), pp. 644–59; Franco Modigliani and Charles Steindel, "Is a Tax Rebate an Effective Tool for Stabilization Policy?" *BPEA*, 1:1977, pp. 175–203; Alan S. Blinder, "Temporary Income Taxes and Consumer Spending," *Journal of Political Economy*, vol. 89 (February 1981), pp. 26–53.

Table 1. Cyclically Adjusted Net Saving as a Percentage of Net National Product, 1971–84^a

<i>Sector</i>	<i>1971–80 average</i>	<i>1980</i>	<i>1981</i>	<i>1982</i>	<i>1983</i>	<i>1984</i>
1. Government	-1.7	-1.4	-1.2	-2.3	-3.4	-4.6
2. Personal	5.5	4.5	5.5	5.2	5.0	5.9
3. Business (net)	2.5	2.1	1.9	2.3	2.9	2.7
4. Total private (2+3)	8.0	6.6	7.4	7.5	7.8	8.6
5. National (1+2+3)	6.3	5.2	6.2	5.2	4.4	4.0

a. Cyclical adjustment is based on regressions of particular saving rates on time, the unemployment rate, and the change in the unemployment rate. The coefficients of the unemployment rate that are used to do cyclical adjustment are as follows (with *t*-statistics in parentheses):

<i>Item</i>	<i>Government</i>	<i>Personal</i>	<i>Business</i>	<i>Private</i>	<i>National</i>
Level of unemployment rate	-0.76 (-5.2)	-0.25 (-2.4)	-0.08 (-0.8)	-0.33 (-2.4)	-1.09 (-8.5)
Change in unemployment rate	-0.17 (-1.1)	0.37 (3.4)	-0.46 (-4.1)	-0.09 (-0.6)	-0.26 (-1.9)

persons (households), businesses, and government on average for the period 1971–80 and then annually for 1980–84.¹⁰ These data do not suggest that the saving rate dropped after the 1981 tax act was passed. As cyclically adjusted government dissaving rose steadily from 1.2 percent of net national product (NNP) in 1981 to 4.6 percent of NNP in 1984, the cyclically adjusted personal saving rate did fall slightly in 1982 and 1983. But the adjusted rate of net business saving rose by more, so that, if households “see through the corporate veil” by treating the retained earnings of corporations as their own, the relevant saving concept (total private saving) actually increased slightly in 1982 and 1983. But a simple look at the data is not the proper way to answer the question. We need to study whether consumers’ reactions to the Reagan “temporary tax increases” were consistent with their behavior in 1968 and 1975.

Interest Rates and Inflation. Recent policy initiatives have focused attention anew on the sensitivity of saving to the after-tax rate of return. Until a decade ago, the relatively small variance in the after-tax real interest rate made inferences about the interest elasticity of consumption

10. Cyclical adjustment was performed by first estimating annual ordinary least squares (OLS) regressions of the form: $s_t = a + bt + cU_t + d[U_t - U_{t-1}] + e_t$ over the period 1954–84. Here s is any of the net saving rates listed in the table, U is the civilian unemployment rate, and t is time. The estimates of the coefficients c and d were used to compute cyclically adjusted saving rates, that is, the saving rates that would have occurred if actual U had been equal to Robert Gordon’s estimate of the natural rate of unemployment each year. See Robert J. Gordon, *Macroeconomics* (Little, Brown, 1983).

tenuous at best.¹¹ Recent years, with their unprecedentedly high and volatile real after-tax interest rates, have remedied that problem and should permit sharper inferences both about interest sensitivity and about any direct effects that inflation might have on consumption. In particular, several years ago, Deaton suggested that the apparent depressing effect of inflation on consumption might be due to the fact that shoppers mistake nominal price increases for real price increases when inflation is unanticipated.¹²

Budget Deficits. The Barro hypothesis, which says (roughly) that government spending should replace taxes in the definition of disposable income, was also difficult to test until recently because cyclically adjusted budget deficits were small and varied little, except during wars. Recent events have changed that. If Barro is right, private saving—and probably personal saving—should have risen dramatically to counteract the effects of government dissaving. Instead, table 1 shows that the cyclically adjusted personal saving rate rose by only 0.4 of a percentage point between 1981 and 1984, while the cyclically adjusted government dissaving rate rose by 3.4 percentage points. If we look more broadly at total private saving,¹³ the rise in the cyclically adjusted saving rate is still only 1.2 percentage points—about one-third of the decline in government saving. Thus, on quick inspection, the data appear hostile to Barro's hypothesis: as the government deficit increased, net national saving fell from 6.2 percent of NNP (close to the average of the previous decade) in 1981 to only 4 percent of NNP (the lowest level in the 1954–84 period) in 1984. But, once again, a more serious econometric investigation is in order.

Plan of the Paper

It has become traditional to divide aggregate consumer spending into two components: purchases of nondurable goods and services, C_t , and

11. Eugene Fama, "Short-Term Interest Rates as Predictors of Inflation," *American Economic Review*, vol. 65 (June 1975), pp. 269–82.

12. See Angus S. Deaton, "Involuntary Saving through Unanticipated Inflation," *American Economic Review*, vol. 67 (December 1977), pp. 899–910.

13. The equivalence hypothesis, it seems to us, suggests that a decrease in government saving should be offset by an increase in personal saving. But if households not only see through the corporate veil but actually reach through it and get corporations to do their bidding, then business saving could offset the government's actions. That is why we look also at total private saving.

purchases of durables. Modern research on the consumption function has focused on the former and typically has modeled C_t in isolation from purchases of durables. An alternative procedure is to work with the sum of the flow of services from durable goods and purchases of nondurable goods and services. We adopt instead a middle position in which each component is treated separately, but relative price effects potentially matter, as they would in a system of demand equations. To facilitate comparison with the recent literature, and because modeling expenditures on durables presents a host of special problems, the present paper deals only with nondurable goods and services.

We begin by developing a baseline consumption function for C_t that includes as arguments such standard variables as income, wealth, interest rates, relative prices, and inflation. We use this function to address a variety of basic questions. Is there any point in decomposing income changes and other variables into anticipated and unanticipated components, or is the traditional specification that ignores this distinction adequate? Is it only surprise changes in variables like income and wealth that matter? Is consumption sensitive either to interest rates or to inflation? Are there detectable relative price effects? Has consumer behavior changed during the Reagan years?

Next, we consider various ways to augment the baseline specification in order to test some more controversial hypotheses. Are permanent and temporary tax changes treated differently by consumers? Is the Barro equivalence hypothesis supported? The last section summarizes what we think we have learned.

A Basic Consumption Function

THE DATA

For purposes of this study, which covers the period 1954:1 to 1984:4, we make four changes in the official national income and product accounts (NIPA) data.

First, we remove the 1975 tax rebate from the income data until we are ready to deal explicitly with the temporary tax issue. We do this so as not to allow the 1975:2 observation to exert undue influence on our baseline specification, for the raw data show a stunning 23.8 percent

annual growth rate of real disposable income in 1975:2, followed by a 5.7 percent annual rate of decline in 1975:3.

Second, the NIPA include gross interest payments from businesses to individuals (for example, on corporate bonds) in personal income, and hence in disposable income, without netting out interest payments from individuals to businesses (for example, for consumer credit). But the tax system essentially nets one against the other and levies taxes only on net interest received. So we subtracted interest paid by consumers to businesses from the NIPA definition of disposable income. This change lowers Y slightly without affecting C .¹⁴

Third, in the NIPA, personal “nontax payments” are grouped with personal taxes. A closer examination of this category reveals that such things as tuition payments to state colleges and fees collected by government hospitals are part of state and local “nontax payments.” Because these items seem more accurately classified as personal consumption, not as taxes, we adjusted the national accounts by subtracting state and local nontax payments, which totaled \$46 billion in 1984, from taxes and from government purchases and adding them to consumption. In doing so, we deflated state and local nontax payments by the NIPA deflator for consumption of services. This adjustment raises C and Y equally without changing the government deficit.

Fourth, since expenditures on durable goods and expenditures on nondurables and services almost certainly require different econometric explanations, they also require some classification scheme. We followed the official U.S. Bureau of Economic Analysis (BEA) classification with one exception—we reclassified clothing and shoes, which are nondurable goods according to NIPA, as durables.

The resulting series on real purchases of nondurable goods and services, when put on a per capita basis, became the basic variable to be studied and is henceforth denoted by C .¹⁵

14. The NIPA count interest paid as part of “personal outlays,” but not as part of “personal consumption expenditures.” The adjustment we make is not a big one. For example, in 1984 interest paid amounted to \$78 billion, while disposable income was \$2,577 billion.

15. All series are seasonally adjusted. We used total population in making per capita conversions. Our implied deflator for total consumer spending (henceforth P), which we use to deflate our version of disposable income, differs trivially from the NIPA deflator because the two definitions of total consumer spending differ slightly.

FUNCTIONAL FORM AND ESTIMATION ISSUES

While the permanent income model perhaps leads more naturally to a linear relationship between consumption and some concept of income, we adopt a logarithmic specification here. We do not believe that any important results are sensitive to the choice between linear and logarithmic form. But since we are interested in studying several price-like variables, such as interest rates, the logarithmic form is convenient in that it avoids the need for numerous interaction terms to allow all the coefficients to depend, for example, on the interest rate. In addition, the modern Euler-equation or "surprise" consumption functions that we wish to study predict that, in the absence of new information, consumption grows from period to period at a rate that depends on the real rate of interest. Once again, this is most easily modeled using a logarithmic specification.

We do not, however, work with the currently fashionable "first differences only" specification, because differencing obliterates the single most striking characteristic of the time series—the remarkable constancy of the C/Y ratio—and therefore is silent about the steady-state properties of the system. Instead we attempt to capture both the short-run dynamics and the long-run properties of the consumption-income relationship by adopting a flexible distributed lag model that accommodates, or "nests," many of the specifications that have been discussed in the literature—including both "Euler-type" specifications and the error-correction model that has been much recommended by David Hendry and several collaborators in the United Kingdom.¹⁶

Specifically, our basic functional form is:

$$(1) \quad \Delta c_t = \beta_0 + \beta_1 c_{t-1} + \beta_2 y_t + \beta_3 y_{t-1} + q_t \delta + z_{t-1} \gamma + u_t,$$

where c and y denote the natural logarithms of consumption and income. In addition to income, there are two types of right-hand variables in equation 1. The q variables are contemporaneously dated variables like

16. James E.H. Davidson, David F. Hendry, Frank Srba, and Stephen Yeo, "Econometric Modelling of the Aggregate Time-Series Relationship between Consumers' Expenditure and Income in the U.K.," *Economic Journal*, vol. 88 (December 1978), pp. 661-92.

wealth, inflation, and relative prices. The z_{t-1} variables are either lagged values of q variables or other variables that are known at time $t-1$ (such as a time trend). The list of z and q variables changes as we examine various hypotheses during the course of the paper.

Whenever we estimate a version of equation 1, we also estimate an augmented specification that decomposes y and q into anticipated and unanticipated components, namely:

$$(2) \quad \Delta c_t = \beta_0 + \beta_1 c_{t-1} + \beta_2 E y_t + \beta_2^* (y_t - E y_t) \\ + E q_t \delta + (q_t - E q_t) \delta^* + z_{t-1} \gamma + u_t.$$

We do this for two reasons. The first is to test the “surprises only” prediction of the pure PIH without liquidity constraints, that is, the implication that changes in consumption from $t-1$ to t should be independent of information about income and wealth that was available at time $t-1$. We perform this test by dropping lagged consumption from the right-hand side of equation 2, re-estimating the equation, and then testing whether the coefficients on anticipated and lagged income and wealth are zero.¹⁷ This test is precisely the “excess sensitivity” test carried out by Marjorie Flavin and others, adapted to a more elaborate specification.¹⁸

The second reason for decomposing all contemporaneous variables into anticipated and unanticipated components is econometric. The current value of a q variable (for example, the relative price of durable goods) might have a well-defined theoretical role in the consumption function. But in addition, the “news” contained in any contemporaneously dated variable might induce consumers to revise their estimates of permanent income, and thus to change their consumption. If so, these variables will be correlated with the change in consumption for reasons

17. This test is sensitive to assumptions about the presence of transitory consumption. If transitory consumption is important, the differenced equation will have moving average residuals, in which case inclusion of lagged c will yield inconsistent estimates, and even without lagged c , the serial correlation will invalidate the standard test statistics. However, we found no evidence of serial correlation whether or not lagged c was included (which is itself evidence against the pure PIH), so we have no reason to doubt the validity of our test statistics. Since c_{t-1} appears to be important in several of the regressions, we suppress it only to make our tests comparable with those in the literature; test statistics when c_{t-1} is included are much less favorable to the surprises-only hypothesis.

18. See Flavin, “The Adjustment of Consumption.”

having nothing to do with their inherent roles in the consumption function, such as the substitution effects of relative prices. The problem is solved by constructing instrumental variables for the q variables from first-stage regressions using data dated $t-1$ and earlier. This amounts to replacing each q by its "anticipated" value. We then allow separately for the effects of the unanticipated q variables by including them in the regressions as well.

Estimation of equation 2 requires a way to deal with unobserved expectations. As just suggested, we adopt the now standard method of generating expectations as the one-period-ahead forecasts from an estimated vector autoregression (VAR) for the variables y and q :

$$(3) \quad \begin{pmatrix} y \\ q \end{pmatrix} = AV_{t-1} + e_t.$$

Here the vector V_{t-1} includes two lags each of c , y , and every variable in q , plus all the z_{t-1} variables and a quadratic time trend. The expected component is the predicted value. The unanticipated component is the series e_t .

The VAR equations themselves are of limited interest and change every time we alter the specification of z and q ; hence they are not reported. However, since the unanticipated, or "surprise," variables generated by these equations play a major role in our analysis, a little description is in order. The VAR equations fit the data quite well, so that most of the variance of y and q is classified as "anticipated" (for example, R^2 for income exceeds 0.99). All the variables are strongly autoregressive (generally of second order), and, in addition, the stock of durable goods helps predict both income and inflation, the nominal rate of interest helps predict wealth, and time helps predict inflation. The appendix reports the "data" on anticipated and unanticipated changes of income, wealth, and inflation generated by one important version of the model. The simple correlations between the actual changes and the surprises are 0.76, 0.66, and 0.74 for income, wealth, and inflation, respectively.

After the VARs are estimated by ordinary least squares (OLS) and used to create anticipated and unanticipated series, equation 2 is estimated by OLS. This simple two-step procedure has much to recommend it over more complicated one-step procedures that treat equation 2 and equation 3 as a system. First, it is simpler computationally. Second, the estimated coefficients in equation 2, the equation of interest, are less

contaminated by specification errors in auxiliary equation 3.¹⁹ However, the two-step procedure does not yield correct estimates of all the standard errors, because it treats the anticipated and unanticipated variables as known data, rather than as the statistical estimates that they are. The standard errors for the coefficients of surprise variables are correct as calculated by OLS. But standard errors for the other coefficients must be obtained from a two-stage least squares (2SLS) regression that omits the surprise terms and uses the VAR as the first stage.²⁰

In applied econometrics, as in life, you rarely get something for nothing. Some assumptions must be made in order to identify a system like equation 2 and equation 3. Our main identifying, and thus untestable, assumption is that transitory consumption, that is, the disturbance in equation 2, is orthogonal to the surprises in income and in other variables, that is, to the disturbances, e_t , in equation 3. While this assumption can certainly be questioned, it is far weaker than Friedman's original assumption that transitory income and transitory consumption are uncorrelated—because an income surprise does lead directly to a consumption surprise according to equation 2.²¹

Note also that the estimated coefficients of the no-surprise variables in equation 2 are precisely those that would be produced by estimating equation 1 by 2SLS. Put differently, the surprise model (equation 2) is observationally equivalent to the traditional model (equation 1) with simultaneity affecting y and q . A model with both simultaneity and surprises is therefore not identifiable; the reader, like the authors, must choose one interpretation or the other.²² Throughout the paper, we adopt the surprise interpretation. But readers preferring the simultaneity interpretation can disregard the coefficients of the surprise variables

19. This is just the standard argument for favoring limited-information over full-information methods.

20. See Adrian Pagan, "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review*, vol. 25 (February 1984), pp. 221–48.

21. Flavin, in "The Adjustment of Consumption," identifies the model instead by assuming a value for the coefficient of unanticipated income in equation 2, a procedure we do not find appealing.

22. Robert Hall, in "The Role of Consumption in Economic Fluctuations," in Robert J. Gordon, ed., *The American Business Cycle* (University of Chicago Press, forthcoming), explores the possibility of estimating (a very parsimonious) consumption function with minimal use of exogeneity assumptions. See also the comments by Deaton that follow Hall's paper.

in equation 2 and treat the other coefficients as 2SLS estimates of equation 1.

For each pair of regressions, we routinely carried out a set of diagnostic tests. First, in addition to the Durbin-Watson statistic, which is biased against finding serial correlation owing to the presence of the lagged dependent variable, and the Box-Pierce Q statistic, we constructed a version of the Lagrange multiplier test for serial correlation up to order four. In none of the regressions was there ever the slightest hint of serial correlation, so we refrain from reporting all these test statistics. This finding is of some interest, however, since the pure PIH implies that the error term should be serially correlated (see footnote 17). Second, we report the results of Chow stability tests over the two halves of the sample and across the Reagan subperiod, 1981:3 to 1984:4; marginal significance levels for rejecting parameter stability are labeled “half-sample stability” and “Reagan stability” in the tables. Except in a few cases to be noted below, parameter stability could not be rejected.

In addition to these diagnostics, we report tests of two more economically interesting hypotheses. First, we always test the hypothesis that the long-run elasticity of consumption with respect to income is unity—something that is not obviously contradicted by the raw data. Second, in all regressions that include the rate of interest, we test the hypothesis that only the real rate of interest matters, that is, that the nominal after-tax interest rate and expected inflation (or actual inflation in the no-surprise regressions) have equal and opposite coefficients.

A final set of tests pertains to the validity of the pure PIH and to the value of decomposing variables into anticipated and surprise components. The hypothesis called “no decomposition” in the tables is that the coefficients of all anticipated and unanticipated variables are equal, so that only actual variables matter, that is, the hypothesis that equation 2 can be reduced to equation 1. Since the PIH suggests so strongly that anticipated and unanticipated income and wealth should get different coefficients, we next test the weaker hypothesis that the decomposition is irrelevant only for these two variables, leaving other variables unconstrained.²³ Finally, we test the surprises-only hypothesis that was

23. Robert Hall pointed out at the September 1985 meeting of the Brookings Panel that, under some circumstances, the pure PIH is fully consistent with our no-surprise specification. Specifically, if income follows a first-order autoregressive process, then y_t is the only variable relevant to predicting future y variables. The appearance of lagged

described earlier, the hypothesis that lagged and anticipated values of income and wealth can be excluded from the regression. The reader is reminded that this test is not based on the surprise regressions reported in the tables, but rather on a constrained version that omits c_{t-1} from the right-hand side (see footnote 17 above).

SIMPLEST SPECIFICATION: THE CONSUMPTION-INCOME RELATIONSHIP

We warm up by estimating the simplest possible consumption function: equations 2 and 1 with no z or q variables. Results are shown as regressions 2.1 and 2.2 in table 2. In this simple specification, which we estimate only because consumption-income relationships like this appear so often in the literature, the surprise and no-surprise versions are very similar, and we cannot come close to rejecting the hypothesis that the two income coefficients in regression 2.1 are equal. By contrast, the hypothesis that only surprises matter is rejected at about the 3 percent level. The implied steady-state elasticity of C to Y is 0.88, which is different from 1.0 at the 10 percent (but not the 5 percent) level. Neither longer lags of c nor longer lags of y were significant when added to regression 2.1 or 2.2.

This specification is deficient in many respects. As a partial remedy, we make three changes in moving to regressions 2.3 and 2.4 in table 2. First, we add wealth to the specification by including three new variables: the logarithm of real wealth, w_t , divided into anticipated and unanticipated components in regression 2.3, and the lagged value of this variable. Our measure of wealth is the household net worth variable used in the MPS model, except that we adjust the value of the government debt from par to market.²⁴

values of c and y in equation 1 can then be rationalized by assuming that transitory consumption is a first-order autoregressive process. Hall is correct that it is very difficult to discriminate between partial adjustment, which we tacitly assume, and serial correlation in the disturbance. However, the data strongly reject the idea that income is a first-order autoregressive process. And neither theory nor the empirical results suggest that the error in the consumption function is serially correlated. So we prefer our interpretation over Hall's.

24. The data necessary to do this come from W. Michael Cox, "The Behavior of Treasury Securities: Monthly, 1942-1984" (Federal Reserve Bank of Dallas, February 1985), and were kindly provided by him.

Table 2. Simple Consumption Functions, 1954:1-1984:4^a

Independent variable, summary statistic, and test of hypotheses	Regression			
	With total income		With labor income	
	Surprise 2.1	No-surprise 2.2	Surprise 2.3	No-surprise 2.4
Constant	0.037 (1.9)	0.036 (2.0)	0.63 (2.6)	0.58 (2.5)
Time trend ($\times 10^3$)	0.50 (2.8)	0.47 (2.8)
Lagged consumption	-0.052(-1.6)	-0.049(-1.7)	-0.188(-3.7)	-0.190(-3.8)
Income				
Lagged	-0.268(-2.8)	-0.254(-5.3)	-0.153(-1.7)	-0.143(-3.0)
Anticipated	0.313 (3.0)		0.236 (2.6)	
Unanticipated	0.294 (5.7)	0.298 (6.4)	0.228 (4.7)	0.228 (5.4)
Wealth				
Lagged	-0.034(-0.8)	-0.068(-2.4)
Anticipated	0.050 (1.0)	0.088 (3.1)
Unanticipated	0.119 (3.1)	

Summary statistic

R^2	0.260	0.260	0.399	0.390
Standard error ($\times 100$)	0.407	0.405	0.373	0.373
Sum of square residuals ($\times 100$)	0.197	0.197	0.160	0.163
Degrees of freedom	119	120	115	117
Long-run income elasticity	0.88	0.88	0.53	0.56

Test of hypotheses

(marginal significance levels)^b

Unit elasticity	0.070	0.059	0.004	0.007
Surprises only	0.032	...	0.038	...
No decomposition	0.87	...	0.46	...
Half-sample stability	0.17	0.10	0.45	0.25
Reagan stability	0.56	0.55	0.66	0.72

Sources: Equations 1 and 2, with data from national income accounts.

a. Quarterly data. The dependent variable is the change in log of total consumption of nondurable goods and services. For each variable, x_t , the anticipated variable is defined as Ex_t , the unanticipated variable as $x_t - Ex_t$. Variables are expressed in natural logarithms (except time). Numbers in parentheses are t -statistics.

b. The "marginal significance level" is the probability, expressed in a decimal, of getting the indicated coefficient if the hypothesis is correct. The hypotheses are as follows. Unit elasticity: the long-run elasticity of consumption to income, assuming that wealth is proportional to income in the long run, is 1.0; no decomposition: equal coefficients of all anticipated and unanticipated variables, so the surprise regression reduces to the no-surprise regression; surprises only: the coefficients of lagged and anticipated income and wealth are zero in a version of the surprise equation that omits c_{t-1} from the right-hand side; half-sample stability: equal coefficients in the two halves of the sample; Reagan stability: equal coefficients before and after 1981:3.

Second, in recognition of the fact that C is not total consumer spending, we add a time trend to the specification. In doing this, we recognize the danger of picking up a spurious time trend.²⁵ Nonetheless, a time trend seems called for on economic grounds; it could represent slowly evolving changes in tastes between durables and nondurables, technological or other changes in the available menu of goods, or other things.

Third, the basic PIH-LCH theory suggests that consumption depends on current wealth and on current and expected future *labor* income. Income from capital is omitted because the current market value of wealth is the best estimate of the discounted stream of income that will be derived from the assets currently owned. The problem, of course, is that disposable labor income is difficult to measure because the income tax is based on total income, not on labor and capital income separately. We constructed a real per capita disposable labor income series, YL , as follows. Starting with the NIPA breakdown of personal income into its components, we apportioned proprietors' income into labor and capital components in the ratio L/I , where:

L is the sum of wages and salaries plus other labor income, and

I is the sum of interest, dividends, and rental income.

Then we attributed personal income taxes to labor and capital in the same ratio, treated all contributions for social insurance as deductions from labor income, and attributed all transfer payments to labor income.

Once these three changes are made, the fits of the equations improve tremendously; R^2 rises by about 0.13, and the standard error falls by about 8 percent. There is no indication of parameter instability. Believers in life-cycle theory will find comfort in the fact that, when regression 2.4 is run using total income rather than just labor income, thereby including both wealth and the income from wealth in the formulation, a non-nested hypothesis test unambiguously selects labor income over total income as the appropriate income variable.²⁶

25. The reasons are discussed in N. Gregory Mankiw and Matthew Shapiro, "Trends, Random Walks, and Tests of the Permanent Income Hypothesis" (Yale University, 1984), and Angus S. Deaton, "Life-Cycle Models of Consumption: Is the Evidence Consistent with the Theory?" (Princeton University, 1985).

26. The test is the Cox test described in M.H. Pesaran, "On the General Problem of Model Selection," *Review of Economic Studies*, vol. 41 (April 1974), pp. 153-71. The test rejects the hypothesis that total income is the correct variable, but cannot reject labor income. Using labor income, not total income, as the appropriate income variable makes irrelevant the common argument that the part of interest income (and expense) that

The wealth coefficients in regression 2.3 show that the significant wealth effect in regression 2.4 derives from unanticipated changes in wealth, as suggested by life-cycle theory. This finding will persist throughout the study. Because each specification has a sizable upward time trend, the estimated long-run elasticity of consumption to income falls to below 0.6, and a unitary elasticity can be rejected decisively.²⁷ Since the impact elasticities are about 0.23, the dynamics work out faster than they did in the previous regressions.

When income is redefined, the income coefficients change moderately (compare regression 2.1 with regression 2.3, or 2.2 with 2.4), the coefficients of expected and unexpected income in regression 2.3 remain very close to one another, and expected income remains quite significant. On the basis of a comparison of regressions 2.4 and 2.3, the data cannot reject the “no decomposition” hypothesis that only actual values of y and w matter (marginal significance level = 46 percent). But the opposite extreme hypothesis that only surprises matter is still rejected (marginal significance level = 3.8 percent).

RELATIVE PRICE TERMS

Intertemporal Prices. All the equations in table 2 omit several potentially important variables, such as inflation and interest rates. In theory, the after-tax real interest rate should influence intertemporal choice; that, after all, is the basic insight of Euler equations. But what is the effective tax rate on interest income, and how should expected inflation be measured?

We measure one minus the effective marginal tax rate on interest $(1 - \tau)$ as the ratio of the yield on (tax-exempt) five-year AAA municipal

represents compensation for inflation should be deducted from standard measures of income in arriving at a true concept of Hicksian income. However, the behavior of liquidity constrained consumers is governed by cash flow, not by Hicksian income. And, for them, the need to make interest payments, even if the payments merely compensate for inflation, might deter consumption. We tested for this by adding interest payments to our regressions, but found erratic and insignificant effects.

27. Throughout the paper, we calculate the long-run income elasticity keeping the wealth-income ratio constant, so that both direct income and indirect wealth effects are included. If the time trend is omitted, the estimated long-run elasticity is very close to 1.0. It is hardly surprising that the presence or absence of a time trend has a dramatic effect on the estimated long-run elasticity.

bonds to the yield on (taxable) five-year AAA corporate bonds. In the sample, τ varies between 0.18 and 0.39. We then construct an after-tax nominal rate of interest by multiplying the three-month Treasury bill rate by $(1 - \tau)$.²⁸ The regressions use the interest rate known at time $t - 1$ for carrying wealth forward to time t because this is the natural timing to use if we think of our specification as representing the intertemporal choice between consumption in periods $t - 1$ and t .²⁹

Inflationary expectations are generated from our VAR, based on the deflator for total consumer spending, called P . Both anticipated inflation ($E\Delta p$) and unanticipated inflation ($\Delta p - E\Delta p$) are entered into regressions 3.1 and 3.2 in table 3, the former as a natural companion to the nominal interest rate, the latter to represent the price confusion effect proposed by Deaton. As noted earlier, however, we tested rather than imposed the constraint that the coefficients of the nominal interest rate and expected inflation be equal and opposite—as would be true if only the real interest rate mattered.

The regression results suggest that nominal interest rates are of little, if any, importance for consumption. Comparing regressions 3.1 and 3.2 shows that the coefficient of r_{t-1} is much larger in absolute value in the surprise version, that is, when contemporaneous variables are instrumented. This was true in every specification.³⁰

Inflation, on the other hand, is quite significant in both specifications, though only in unanticipated form in regression 3.1. Note, however, that expected inflation gets the wrong sign in regression 3.1. If only real interest rates matter, $E\Delta p$ and r should have equal and opposite coefficients. Though this hypothesis cannot be rejected at the 10 percent level owing to large standard errors, the sum of the coefficients is far from zero. Why it should be nominal interest rates, not real rates, that matter is a puzzle that has cropped up in other contexts.³¹ We have no good

28. As an alternative, we also tried using the five-year municipal bond rate. Results were similar. Since the short rate makes better theoretical sense, and usually provided a slightly better fit, we report only those results.

29. In response to suggestions from some readers, we also tried adding current r to the regression. See footnote 30.

30. When r_t was added to the regressions, the results did not change much. Specifically, r_t and r_{t-1} generally got large coefficients with opposite signs, reflecting collinearity between the two variables. The sum of the two coefficients was small and negative.

31. See, for example, Christopher Sims, "Comparison of Interwar and Postwar Business Cycles: Monetarism Reconsidered," *American Economic Review*, vol. 70 (May 1980), *Papers and Proceedings*, 1979), pp. 250–57; Robert Litterman and Laurence Weiss,

explanation to offer. But, in any case, the best guess at this stage is that the significant coefficient of inflation in regression 3.2 reflects the price confusion effect.

There are no notable changes in coefficients when interest rates and inflation are added to the regressions, although the hypothesis that only surprises matter can no longer be rejected.

Relative Prices of Different Goods. In principle, consumers should respond not only to the relative price of the same good in different time periods, but also to the relative prices of different goods in the same time period. This is easy to test.

From among the prices of the three main components of consumer spending—*PS*, the price of services; *PND*, the price of nondurable goods; and *PD*, the price of durables—we selected the latter two for the construction of relative price variables. Regression 3.4, the no-surprise specification, adds the two relative prices in both current and lagged form; in logs, the new variables are $(pn - p)_t$ and $(pn - p)_{t-1}$ for nondurables and $(pd - p)_t$ and $(pd - p)_{t-1}$ for durables. However, when we tried to enter anticipated, unanticipated, and lagged relative prices in the surprise specification, regression 3.3, the anticipated and lagged variables were almost perfectly correlated. Consequently, 3.3 omits the lagged relative prices.

In addition, our regressions ought to include any other variable that has a strong influence on the allocation of income between durable goods and *C*. The stock of durables at the start of the quarter, K_t , is an obvious candidate, since, given the desired stock, a higher opening stock ought to lead to lower spending on durables, and hence to higher spending on *C*, other things being equal.

Regressions 3.3 and 3.4 show a substantial improvement in fit over their predecessors— R^2 rises by about 0.08, and the standard error of the equation falls by about 5 percent. Five of the eight relative price variables in the two regressions are significant. In particular, the large negative coefficients of $(pn - p)_t$ in regression 3.4 and of the surprise in $(pn - p)_t$ in regression 3.3 both suggest a strong transient effect of the relative price of nondurable goods. In either specification, a 1 percent rise in

“Money, Real Interest Rates, and Output: A Reinterpretation of Postwar U.S. Data” (Federal Reserve Bank of Minneapolis, January 1984); and Alan S. Blinder, “Retail Inventory Behavior and Business Fluctuations,” *BPEA*, 2:1981, pp. 443–505.

Table 3. Consumption Functions with Relative Prices, 1954:1-1984:4^a

<i>Independent variable, summary statistic, and test of hypotheses</i>	<i>Regression</i>			
	<i>Intertemporal prices only</i>		<i>All relative prices</i>	
	<i>Surprise</i> 3.1	<i>No-surprise</i> 3.2	<i>Surprise</i> 3.3	<i>No-surprise</i> 3.4
Constant	0.47 (1.5)	0.39 (1.6)	-0.25 (-0.6)	0.17 (0.6)
Time trend ($\times 10^3$)	0.42 (2.0)	0.37 (2.1)	0.19 (0.8)	0.34 (1.9)
Lagged consumption	-0.154(-2.5)	-0.172(-3.2)	-0.099(-1.3)	-0.176(-3.0)
Income				
Lagged	-0.098(-0.7)	-0.080(-1.5)	0.074 (0.5)	-0.071(-1.4)
Anticipated	0.169 (1.1)		0.003 (0.02)	
Unanticipated	0.170 (3.3)	0.172 (3.5)	0.166 (3.2)	0.164 (3.4)
Wealth				
Lagged	-0.004(-0.1)	-0.051(-1.8)	0.047 (0.9)	-0.012(-0.4)
Anticipated	0.019 (0.4)		-0.005(-0.1)	
Unanticipated	0.098 (2.4)	0.072 (2.4)	0.118 (3.0)	0.056 (1.9)
After-tax nominal interest rate, lagged	-0.190(-0.9)	-0.016(-0.1)	-0.231(-1.1)	-0.139(-1.0)
Inflation				
Anticipated	-0.117(-0.6)		-0.881(-2.9)	-0.261(-1.8)
Unanticipated	-0.332(-2.6)	-0.238(-2.5)	-0.096(-0.6)	
Relative price of nondurable goods				
Lagged	0.270 (2.7)
Anticipated	0.127 (2.7)	-0.216(-2.0)
Unanticipated	-0.247(-2.2)	

Relative price of durable goods			
Lagged	0.030 (0.3)
Anticipated	...	0.078 (2.2)	0.009 (0.1)
Unanticipated	...	0.031 (0.3)	
Stock of durable goods			
Anticipated	...	0.001 (0.02)	0.003 (0.2)
Unanticipated	...	0.153 (0.8)	
<i>Summary statistic</i>			
R^2	0.443	0.422	0.497
Standard error ($\times 100$)	0.364	0.366	0.349
Sum of square residuals ($\times 100$)	0.149	0.154	0.134
Degrees of freedom	112	115	110
Long-run income elasticity	0.56	0.66	0.79
<i>Test of hypotheses</i>			
<i>(marginal significance levels)</i>			
Unit elasticity	0.11	0.073	0.30
Real interest rate	0.33	0.12	0.04
No decomposition	0.26
No decomposition for income and wealth	0.45
Surprises only	0.85
Half-sample stability	0.34	0.13	0.21
Reagan stability	0.38	0.40	0.42

Sources: Same as table 2.

a. See notes to table 2. The hypothesis called "real interest rate" is that only the real after-tax interest rate, not the nominal rate and expected (or actual) inflation separately, enters the regressions. For each variable, x_t , the anticipated variable is defined as Ex_t , the unanticipated variable as $x_t - Ex_t$.

PND relative to P (only if unanticipated in regression 3.3) reduces the annual growth rate of C by about 0.9 percent in the first quarter. The estimated steady-state (or anticipated) effects of $pn - p$ are positive, however.³²

Results are weaker for the relative price of durables, $pd - p$, although the steady-state effect appears to have the correct (positive) sign in both specifications. While the coefficient of the opening stock of durables gets insignificant coefficients in both specifications, we leave it in the regression for theoretical reasons lest it interact in important ways with any other variable.

Perhaps more interesting than the estimated relative price effects per se is the way the addition of relative prices changes some of the other coefficients. Two changes are particularly notable. First, regression 3.3, unlike regression 3.1, strongly suggests that only unanticipated changes in income cause consumption to change. Second, the conclusion in regression 3.1 that unanticipated inflation matters more than anticipated inflation is dramatically reversed in regression 3.3. In 3.3, and in the rest of the regressions estimated for this study, the effect of inflation on consumption seems to derive from anticipated inflation. The interest elasticity is a bit stronger, though still not significant, when relative prices are included. And the constraint that the coefficients of r and $E\Delta p$ are equal and opposite can now be rejected strongly in regression 3.3 (marginal significance level = 0.5 percent).

Not surprisingly, given these changes in coefficients, we can now reject, at the 0.8 percent level, the hypothesis that the split between anticipated and unanticipated components does not matter. This result nominates regression 3.3 as our best consumption function. In consequence, we display in the appendix the series on unanticipated income, wealth, and inflation that underlie regression 3.3. Nevertheless, we retain both 3.3 and 3.4 as baseline specifications to be used in the next section because regression 3.3, but not 3.4, shows some slight evidence of parameter instability across half samples, and because some economists are, justifiably, skeptical of our method for decomposing variables into

32. The theoretically expected sign for the relative price of nondurable goods is unclear because the consumption measure includes services as well. Note too that the greater importance of the "surprise" as compared with the anticipated relative prices could be interpreted as evidence of simultaneity between consumption and prices. See our discussion above.

anticipated and unanticipated components. Note also that our surprise equation is not the pure surprise version of the PIH, because it includes both lagged and anticipated terms.

Further Investigation of Interest Sensitivity. The implied steady-state semi-elasticity of consumption to the nominal rate of interest in regression 3.3 is -2.3 . That means that, with the path of income held constant, a 1 percentage point rise in r will eventually decrease C by 2.3 percent—a strong effect. The result, however, seems to be fragile. The strong elasticity is to the *nominal* interest rate and does not appear if only the *real* rate is allowed in the regression. Furthermore, the strong negative effect of the nominal interest rate appears only in the surprise version; the corresponding semi-elasticity in the no-surprise regression is only -0.8 . Since there has been so much concern lately with the sensitivity of saving to the rate of return, we decided to look at the interest elasticity issue more deeply.

Regressions 4.1 through 4.4 in table 4 disaggregate C into its two main components, nondurable goods and services, and show, much to our astonishment, that it is actually spending on services, not on goods, that is sensitive to interest rates. We have a hard time believing that this sensitivity represents intertemporal substitution, but cannot offer a better explanation. The “equal and opposite” constraint on the coefficients of nominal interest rates and inflation continues to be rejected at very exacting significance levels for services, but not for nondurable goods.

Dividing services into five component parts (housing, household operation, transportation services, other services, and state and local nontax receipts) shows that the significant negative interest elasticity can be traced mainly to housing and transportation services. It may seem surprising at first that spending on these services should be interest elastic, but each of these service flows relates directly to a durable stock (houses and automobiles), the demand for which is probably quite sensitive to interest rates.³³ That observation does not, of course, explain why it is the nominal interest rate that matters.

33. Several readers suggested that the negative interest elasticity of housing services is an artifact of the BEA's imputation procedures. That does not appear to be the case, however. Services of owner-occupied houses are imputed, of course. But the imputation is based on actual observed rent-to-value ratios on rented houses, deflated by the consumer price index for rent. There is no mechanical linkage between interest rates and imputed housing services.

Table 4. Disaggregated Consumption Functions, 1954:1-1984:4^a

<i>Independent variable, summary statistic, and test of hypotheses</i>	<i>Regression</i>			
	<i>Consumption of nondurable goods</i>		<i>Consumption of services</i>	
	<i>Surprise 4.1</i>	<i>No-surprise 4.2</i>	<i>Surprise 4.3</i>	<i>No-surprise 4.4</i>
Constant	-0.14 (-0.3)	-0.13 (-0.4)	-0.02 (-0.04)	0.33 (1.1)
Time trend ($\times 10^3$)	-0.15 (-0.7)	-0.28 (-1.6)	0.55 (1.2)	0.88 (2.5)
Lagged consumption	-0.254(-3.5)	-0.295(-4.7)	-0.087(-1.0)	-0.161(-2.5)
Income				
Lagged	-0.208(-0.8)	-0.098(-1.1)	0.136 (1.0)	-0.012(-0.2)
Anticipated	0.350 (1.3)		-0.085(-0.6)	
Unanticipated	0.219 (2.6)	0.266 (3.4)	0.091 (1.6)	0.072 (1.4)
Wealth				
Lagged	0.101 (1.2)	0.020 (0.4)	0.042 (0.8)	-0.013(-0.4)
Anticipated	-0.032(-0.4)		-0.017(-0.3)	
Unanticipated	0.161 (2.5)	0.067 (1.4)	0.089 (2.1)	0.039 (1.2)
After-tax nominal interest rate, lagged	0.128 (0.4)	-0.060(-0.3)	-0.420(-2.1)	-0.196(-1.4)
Inflation				
Anticipated	-0.877(-1.6)		-0.570(-2.3)	
Unanticipated	-0.059(-0.2)	-0.194(-0.8)	-0.196(-1.0)	-0.349(-2.1)
Relative own price				
Lagged	...	0.750 (4.6)	...	0.126 (0.7)
Anticipated	0.101 (1.2)	-0.749(-4.2)	-0.125(-2.3)	-0.235(-1.3)
Unanticipated	-0.812(-4.3)		-0.214(-1.1)	

Relative price of durable goods					
Lagged	...	0.261 (1.7)	...	-0.066(-0.9)	
Anticipated	0.112 (1.8)		0.009 (0.4)		
Unanticipated	-0.115(-0.9)	-0.196(-1.2)	0.101 (1.2)		0.080 (1.0)
Stock of durable goods					
Anticipated	0.015 (0.5)	0.003 (0.1)	-0.003(-0.2)		0.008 (0.5)
Unanticipated	0.091 (0.3)		-0.003(-0.02)		
<i>Summary statistic</i>					
R^2	0.533	0.516	0.360		0.325
Standard error ($\times 100$)	0.574	0.574	0.363		0.366
Sum of square residuals ($\times 100$)	0.349	0.362	0.140		0.147
Degrees of freedom	106	110	106		110
Long-run income elasticity	0.89	0.87	0.83		0.58
<i>Test of hypotheses</i>					
<i>(marginal significance levels)</i>					
Unit elasticity	0.66	0.35	0.81		0.13
Real interest rate	0.27	0.42	0.006		0.013
No decomposition	0.0002	...	0.35		...
No decomposition for income and wealth	0.19	...	0.13		...
Surprises only	0.74	...	0.58		...
Half-sample stability	0.26	0.40	0.039		0.019
Reagan stability	0.77	0.24	0.17		0.32

Sources: Same as table 2.

a. See notes to tables 2 and 3. Dependent variables are the change in log of real per capita consumption of nondurable goods for equations 4.1 and 4.2, and the change in log of real per capita consumption of services for equations 4.3 and 4.4. For each variable, x_t , the anticipated variable is defined as Ex_t , the unanticipated variable as $x_t - Ex_t$.

As we had expected, the negative interest elasticities that we find cannot be detected over a shorter data period that excludes the volatile 1980s, even though the equations pass formal stability tests. When we ran the regressions in tables 3 and 4 over a sample ending in 1979:4, the estimated interest elasticities were either insignificantly negative or positive.

THE BASELINE CONSUMPTION FUNCTION

Hereafter, we shall take equations 3.3 and 3.4 to be our baseline consumption functions, so it is worth dwelling a moment on the steady-state properties of the two. The implied steady-state consumption functions from regressions 3.3 and 3.4, respectively, are:

$$C = 7.85 Y^{0.78} W^{0.42} K^{0.005} \left(\frac{PND}{P}\right)^{1.28} \left(\frac{PD}{P}\right)^{0.79} e^{0.0019t - 2.3r - 8.8\Delta p}$$

and

$$C = 6.76 Y^{0.53} W^{0.25} K^{0.015} \left(\frac{PND}{P}\right)^{0.31} \left(\frac{PD}{P}\right)^{0.22} e^{0.0019t - 0.8r - 1.5\Delta p}.$$

The two steady-state specifications differ mainly because of different coefficients on c_{t-1} in regressions 3.3 and 3.4. This illustrates the difficulty of inferring long-run properties from short time series.

Not surprisingly, most of the biggest residuals occur when c changes the most; but no residual in the entire sample is as large as three standard errors. The biggest comes in 1983:2, when our regression fails to capture the huge (and transitory) acceleration of consumer spending. We also miss the upward "blip" in spending that occurred in 1965:4. A plot of the residuals (not shown) has no evident runs of positive or negative residuals. In particular, there is no tendency for the regressions to overpredict consumption growth (that is, produce negative residuals) after the introduction of several special incentives for saving in the August 1981 tax bill; in regression 3.4, the mean residual for the 13 quarters 1981:4–1984:4 is actually +0.07 percent.

Government Policy and Consumption: Tests of Hypotheses

This section tests a variety of more controversial hypotheses about how consumers react to changes in government policy, such as tempo-

rary taxes and budget deficits. In each case, the baseline consumption function of the previous section serves as a starting point.

REACTIONS TO TEMPORARY INCOME TAX CHANGES

Friedman's permanent income hypothesis suggested long ago that consumers should react to temporary income tax changes less than they do to permanent ones, and Lucas made that point part of his famous critique of econometric policy evaluation. Furthermore, several empirical studies have detected such a difference. In the most comprehensive study of the issue, Blinder estimated that consumers treat a temporary tax change as roughly half ordinary income change, half pure windfall.³⁴

To study this issue, we begin, as Blinder did in his previous paper, by dividing income along lines that differ from the usual permanent-transitory distinction. Specifically, we define "special" income, S_t , as the disposable income, positive or negative, derived from temporary tax changes and "regular" income, R_t , as the rest; that is, $R_t = Y_t - S_t$. The motivation for this dichotomy is that special income is observably "more transitory" than regular income, which is, itself, a blend of permanent and transitory components. Because of this difference, the PIH suggests that S should have a smaller effect on consumption than does R ; that is, the appropriate income variable should be:

$$R + \mu S = Y - (1 - \mu)S = Y \left[1 - (1 - \mu) \frac{S}{Y} \right],$$

where μ is an empirically estimated coefficient between zero and one. Since the log of a sum is not the sum of the logs, we approximate the natural log of this income concept as follows:

$$\ln Y + \ln \left[1 - (1 - \mu) \left(\frac{S}{Y} \right) \right] \approx \ln Y - (1 - \mu) \left(\frac{S}{Y} \right),$$

which is an excellent approximation, since the income involved in temporary tax changes is small, always less than 4 percent of other income. Thus, to estimate μ , we simply need to add a new variable, the ratio S/Y , to our regressions. Since $y = \ln Y$ enters in both current and lagged form, S/Y enters in current and lagged form as well. This gives us

34. See Blinder, "Temporary Income Taxes."

two different estimates of $(\mu - 1)$: an impact effect derived from the ratio of the coefficient of $(S/Y)_t$ to that of y_t , and a long-run effect derived from the sums of the (S/Y) and y coefficients.³⁵ The pure PIH suggests that μ should be near zero, while $\mu = 1$ means that actual measured income is the relevant income concept.

Before looking at the results, we need to explain how we constructed the S_t series. We distinguished four different temporary tax change episodes.

The 1968–70 tax surcharge. A one-year temporary tax surcharge began in 1968:3 and was subsequently extended, at reduced rates, for a second year. The 1968–70 episode was studied by Arthur Okun and William Springer with conflicting results.³⁶ For this study, revised estimates of the revenues from the surtax during the eight quarters of its existence were taken from the May 1978 issue of the *Survey of Current Business*.

The 1975 rebate. A rebate of 1974 taxes was paid mostly in May and June of 1975, with traces trickling into the third and fourth quarters of 1975. This episode was studied by Modigliani and Charles Steindel.³⁷

The other 1975 cuts. The 1975 tax act also reduced bracket rates “temporarily” for one year. Since these rate cuts were subsequently extended in the Revenue Adjustment Act of 1975 and again in the Tax Reform Act of 1976, and eventually became a permanent feature of the tax code, we had to make an arbitrary assumption about when consumers ceased viewing them as temporary. We assumed that consumers shifted linearly from viewing the cuts as 100 percent temporary in 1976:1 to viewing them as 100 percent permanent by 1977:1. All three of the above-mentioned episodes were studied by Blinder;³⁸ but data on both 1975 tax reductions have been revised since then, and were kindly provided to us by the BEA.

The Reagan tax cuts. As mentioned in the introduction, the Reagan tax cuts of 1981–84 have given us another “temporary tax” episode, though one of a different character that may have been treated differently

35. In the surprise version, we sum the coefficients of lagged and anticipated variables to obtain the long-run effect. This procedure is followed throughout.

36. Okun, “The Personal Tax Surcharge”; Springer, “Did the 1968 Surcharge Really Work?”

37. Modigliani and Steindel, “Is a Tax Rebate an Effective Tool?”

38. Blinder, “Temporary Income Taxes.”

by consumers. From a strict PIH viewpoint, one should treat the phased-in tax cuts as a large permanent tax reduction in the middle of 1981:3 coupled with a temporary tax increase that diminished gradually to zero by 1984:1. The series is offered for inspection in table 5 and explained with the aid of the following notation:

A_t is income from the actual Reagan tax cuts in quarter t (mostly BEA data from the April 1985 *Survey of Current Business*);

T_t is income from the theoretical "permanent" tax cut applicable to quarter t (constructed by us);

S_t is income from the theoretical temporary tax hike applicable to quarter t ($A_t - T_t$).

We dealt separately with each of the three reductions in withholding rates and the nonwithheld part of the liabilities.

October 1981 withholding reductions. Though the Reagan tax cuts were enacted in August 1981, the first stage was not effective until October. We therefore set T_t for 1981:3 equal to half the BEA value for A_t in 1981:4. Since $A_t = 0$ in 1981:3, $S_t = -T_t$ that quarter. Thereafter, the October 1981 cut was fully effective, so S_t from this source is zero.

July 1982 withholding reductions. BEA data assume that the July 1, 1982, rate reductions became effective on that date. But tax rates apply to calendar years, so half the July 1 cuts actually applied to the entire calendar year. So we began by changing the BEA series on A_t by using quarterly withholding payments to allocate the calendar 1982 tax reduction (\$13.3 billion at annual rates) to the four quarters of 1982. We then estimated T_t for each quarter of 1982 by assuming that the full July 1 rate reductions were effective throughout the year. Having done that, we constructed S_t as $A_t - T_t$. For quarters beyond 1982:4, the full July 1982 rate reductions were in effect, so S_t from this source is zero.

July 1983 withholding reductions. The procedure used here is exactly the same as that for the July 1982 reductions.

Nonwithheld taxes. The BEA provided data on A_t . We constructed T_t by applying the 1984 ratio of nonwithheld tax cuts to withheld tax cuts to the quarters from 1981:4 through 1983:4 (half the rate in 1981:3). Then S_t was $A_t - T_t$ as usual.

There is every reason, however, to treat the S_t series for Reagan's "temporary tax hike" differently from the others because the other three episodes were genuine changes in income flows received by households, while the mythical Reagan "hike" is a theoretical construct based on

Table 5. Estimated Consumer Income from Reagan Tax Reductions, 1981:3–1983:4*
Billions of dollars

<i>Quarter</i>	<i>Actual</i>	<i>Permanent</i>	<i>Temporary</i>
1981:3	0.6	46.0	–45.4
1981:4	18.6	88.0	–69.4
1982:1	39.7	91.1	–51.4
1982:2	40.4	93.4	–53.0
1982:3	40.9	95.4	–54.5
1982:4	41.3	96.8	–55.5
1983:1	81.7	97.0	–15.3
1983:2	83.5	99.4	–15.9
1983:3	85.1	101.7	–16.6
1983:4	87.6	105.0	–17.4

Source: Authors' computations.

a. Starting in 1984:1, "actual" and "permanent" are equal, so "temporary" is zero.

the PIH. The following procedure suggests itself. Suppose the income concept relevant to consumption in 1981:3–1983:4 was:

$$Y_t + a(Y_t^p - Y_t) = Y_t + a(-S_t^*),$$

where $Y_t^p - Y_t = (-S_t^*)$ is the income from promised future tax cuts that is currently withheld; this series is given, with sign reversed, in the last column of table 5. Thus we see that a is analogous, but not necessarily equal, to the $1 - \mu$ parameter for the other three temporary tax episodes; that is, $a = 1$ connotes the pure PIH, while $a = 0$ means that consumers ignored the promised future tax cuts.

Rather than repeat long tables of regression coefficients that hardly change, we report in table 6 and in the tables that follow only the estimates of the new coefficients and parameters of interest, plus the relevant test statistics. We will note any important changes in coefficients as we proceed. When temporary tax change variables were added, the most notable changes were an increase of about 50 percent in the coefficient of lagged consumption and a doubling of the estimated time trend. In consequence, the estimated long-run elasticity of consumption to income fell.

In table 6, the variable S/Y pertains to the 1968–70 and 1975–76 temporary tax changes, while S^*/Y pertains to the 1981–83 temporary tax change. For both episodes, the point estimates suggest that the effects of temporary tax changes are near zero, that is, that they are ignored by consumers. The long-run effects are also near zero for the

Table 6. Tests for Temporary Taxes, 1954:1–1984:4^a

Variable, parameter, and test	Regression	
	Surprise	No-surprise
Ratio of income from temporary tax to total labor income ^b		
1968–70 and 1975–76 episode		
$(S/Y)_t$	-0.152(-1.2)	-0.200(-2.0)
$(S/Y)_{t-1}$	-0.077(-0.7)	-0.120(-1.2)
1981–83 episode		
$(S^*/Y)_t$	-0.014(-0.1)	-0.014(-0.1)
$(S^*/Y)_{t-1}$	-0.010(-0.1)	0.041 (0.3)
<i>Parameter^c</i>		
Long-run income elasticity	0.83	0.62
Estimated μ		
Impact	0.12	-0.06
Impact and lagged	-1.37	-1.51
Estimated a		
Impact	0.08	0.07
Impact and lagged	0.24	-0.22
<i>Test of hypotheses (marginal significance levels)^d</i>		
Strict PIH ($\mu = 0, a = 1$)	0.49	0.059
Measured income ($\mu = 1, a = 0$)	0.50	0.058
Unit elasticity	0.64	0.017
Real interest rate	0.011	0.041
No decomposition	0.074	...
No decomposition for income and wealth	0.24	...
Surprises only	0.98	...

Sources: Same as table 2.

a. The dependent variable is the change in log of consumption of nondurable goods and services. New variables are added to baseline equations 3.3 and 3.4.

b. These are coefficients of the indicated variables in the regressions.

c. As described in the text, the parameters are estimated as the ratio of coefficients from the regression.

d. This is the probability, expressed in a decimal, of getting the indicated coefficients if the hypothesis is correct.

Reagan episode, but are actually strongly negative for the pre-Reagan episodes—which means that consumption moved in the wrong direction. These estimates are quite different from Blinder's 0.50 estimate of a parameter similar to μ published four years ago, based on unrevised data and a very different specification. The source of this discrepancy is an intriguing question for future research.

The null hypothesis $a = 1$ and $\mu = 0$ corresponds roughly to the strict PIH in that it says that the relevant income concept excluded the 1968 and 1975–76 temporary tax changes but treated the Reagan tax cuts as if

they were fully effective in August 1981. This hypothesis can be rejected in the no-surprise version of the model, though only at the 6 percent level; it cannot be rejected at any reasonable level in the surprise version, owing to larger standard errors. The opposite extreme hypothesis, that $a = 0$ and $\mu = 1$, says that consumers spend on the basis of conventional measured income, with no special allowance for temporary tax changes. This hypothesis again can be rejected in the no-surprise version only at the 6 percent level, and not at all in the surprise version.

While these results are probably not precise enough to persuade anyone to abandon strongly held a priori views, they do point toward a mixed conclusion ($a = 0$ and $\mu = 0$) in which consumers do not spend on the basis of income received from temporary tax cuts (in accord with the PIH) but nonetheless ignored the promised tax reductions of 1981–84 until they actually occurred (an anti-PIH result).

Because the 1975 rebate was received so late in the second quarter of 1975, we tried estimating separate coefficients for the rebate on the supposition that $(S/Y)_{t-1}$ might be relatively more important than $(S/Y)_t$ for the rebate than for other temporary tax episodes. This did not turn out to be true, however. We also tried to distinguish between anticipated and unanticipated temporary tax changes. But the estimates were based on so little data that the standard errors were enormous.

IS THE GOVERNMENT VEIL PIERCED?

Our preliminary look at the data in the introductory section made it appear as though households (or households and businesses combined) did not raise their saving to offset the rise in government dissaving after 1981, as would be expected if debt and taxes were equivalent. Here we investigate the Barro equivalence hypothesis more rigorously.

If the Barro hypothesis is right, our consumption function is misspecified in two ways. First, the market value of the government debt should be omitted from household net worth. This is easily done, thanks to extremely accurate estimates of the market value of the debt compiled and kindly provided to us by Michael Cox.³⁹ To test the Barro hypothesis against the traditional specification, it is natural to replace $\ln W$ in our regressions by $\ln(W - D + \mu_1 D)$, where D is the market value of

39. Cox, "The Behavior of Treasury Securities."

government debt. If $\mu_1 = 0$, the Barro hypothesis is correct; if $\mu_1 = 1$, the full government debt is included in wealth (that is, there is no tax discounting). Using the same approximation as before, this amounts to replacing $\ln W$ by two variables: $\ln(W - D)$ and $D/(W - D)$. The ratio of the two coefficients is an estimate of μ_1 .

Second, if Barro is right, disposable labor income is the wrong income concept. But what income concept is right? We answer this question in stages, first supposing that we were working with total income rather than just labor income.

Consumers who pierce the government veil presumably understand that the income actually available to them is net national product (*NNP*)⁴⁰ minus government purchases (*G*) minus transfers and interest paid by the U.S. government to foreigners plus state and local personal nontax payments.⁴¹ Call this concept of income *YB*. Manipulation of some NIPA identities reveals that:

$$(4) \quad YB = Y + RET + SURP,$$

where *Y* is total disposable income as defined by us, *SURP* is the ordinary NIPA government surplus, and the inclusion of *RET*, retained earnings, reflects the hypothesis that the corporate veil is also pierced.

Now we must come to grips with the fact that our baseline specification, regressions 3.3 and 3.4, uses labor income, not total income, and that identity 4 does not hold in labor income. We begin by defining

$$DIFF = YB - Y$$

as the difference between the Barro concept of total income and our own. Since retentions are irrelevant in a specification based on labor income, we eliminate them, which makes $DIFF = SURP$.⁴² Assume that a portion $\lambda_t = L/(L + D)$ of the surplus gets into labor income, where

40. Seater and Mariano, in "New Tests," use GNP rather than NNP. We do not see the argument for including depreciation.

41. Because transfers and interest paid by the U.S. government are sent abroad, they are not available to domestic households. State and local personal nontax payments are considered to be personal consumption, as explained above.

42. To check that this decision did not bias the case against the Barro hypothesis, we ran regressions that included a retentions variable and tested the restrictions that exclude this variable. The restrictions could not be rejected. We also tested the Barro hypothesis on the assumption that retentions should be included as part of labor income. The results were not sensitive to the way retentions were handled.

λ_t is the ratio we used earlier to define labor income. Then a reasonable concept of labor income if the Barro hypothesis is correct is:

$$YBL = YL + \lambda SURP.$$

Hence, using the usual approximation, we can replace $\ln(YL)$ in our regressions by:

$$\ln(YL + \theta\lambda SURP) = \ln YL + \theta\lambda \left(\frac{SURP}{YL} \right).$$

If Barro is right, $\theta = 1$, which means that we should redefine income by adding back labor's share of taxes and subtracting labor's share of government purchases.

But that procedure would saddle Barro's hypothesis with an auxiliary constraint that is no part of the tax discounting idea, namely that government purchases have no direct effect on the marginal utility of private consumption. If, instead, government nondefense purchases, G , are either substitutes or complements for private consumption, there is no reason to enter G and taxes, net of transfers, with equal and opposite signs, as would be done if we used the variable $SURP$. To allow for this, we add to the regression $g = \log G$, both contemporaneously, decomposed into anticipated and unanticipated components in the surprise version, and lagged.⁴³

The results are reported in table 7. The coefficients on the current, or surprise, versions of the new income variable $\lambda_t (SURP/YL)_t$, are very small, but those on the lagged, or anticipated, terms are large enough to suggest that a substantial fraction (55 percent in the no-surprise version and 92 percent in the surprise version) of future taxes is eventually discounted back to the present. The coefficients of the new wealth variables imply estimates of μ_1 of 1.0 or more. Hence, while the point estimates of the income coefficients give some support to the Barro hypothesis, the wealth coefficients are the opposite of what it requires.

This rather odd combination of coefficients probably reflects the fact that the big movements in both of the special Barro variables (for income and for wealth) came at the same time. Unfortunately, because some

43. The coefficients of government spending were generally positive and of marginal significance in the regressions, indicating that government purchases and private consumption are complements rather than substitutes.

Table 7. Tests of Barro Hypothesis, 1954:1–1984:4^a

Variable, parameter, and test	Regression	
	Surprise	No-surprise
Barro income variable ^{b,c}		
Lagged	-0.009 (-0.1)	0.039(1.0)
Anticipated	0.094 (0.7)	0.002(0.05)
Unanticipated	0.008 (0.2)	
Ratio of government debt to other wealth ^c		
Lagged	-0.203 (-0.6)	0.012(0.1)
Anticipated	0.257 (0.6)	0.176(1.3)
Unanticipated	0.157 (0.9)	
<i>Parameter^d</i>		
Long-run income elasticity	1.77	1.23
Estimated θ		
Impact	0.05	0.01
Impact and lagged	0.92	0.55
Estimated μ_1		
Impact	1.18	2.75
Impact and lagged	3.54	3.30
<i>Test of hypotheses</i> <i>(marginal significance levels)^e</i>		
Barro hypothesis ($\theta = 1, \mu_1 = 0$)	0.51	0.016
Traditional hypothesis ($\theta = 0, \mu_1 = 1$)	0.72	0.060
Unit elasticity	0.66	0.63
Real interest rate	0.48	0.53
No decomposition	0.061	...
No decomposition for income and wealth	0.15	...
Surprises only	0.54	...
Half-sample stability	0.033	0.0027
Reagan stability	0.68	0.26

Sources: Same as table 2.

a. The dependent variable is the change in log of consumption of nondurable goods and services. For each variable, x_t , the anticipated variable is defined as Ex_t , the unanticipated variable as $x_t - Ex_t$. New variables are added to baseline equations 3.3 and 3.4.

b. The variable is defined as $\lambda(SURP/YL)$, where λ is the labor share, $SURP$ is the government surplus, and YL is labor income.

c. These are coefficients of the indicated variables in the regressions.

d. As described in the text, the parameters are estimated as the ratio of coefficients from the regression.

e. This is the probability, expressed as a decimal, of getting the indicated coefficients if the hypothesis is correct.

coefficients support Barro neutrality and others contradict it, the equation's implications for the neutrality hypothesis are not transparent. To understand better what the coefficients mean, we ran a simulation in which personal income taxes fell by \$25 per capita (\$100 at annual rates) in 1981:4 only, leaving the government debt higher by \$25 per capita

thereafter,⁴⁴ with all other variables in the regression held constant.⁴⁵ If the neutrality hypothesis were correct, such an event would have no effect on consumption. The actual simulated effects on the level of real per capita consumption of nondurable goods and services are given in table 8, where they are juxtaposed against the results implied by the traditional model ($\theta = 0$, $\mu_1 = 1$). The results are far closer to the implications of the traditional view than they are to the zeros called for by the Barro hypothesis.

In the no-surprise specification, a formal test rejects the Barro hypothesis at the 1.6 percent level, though, with a marginal significance level of 6 percent, the traditional view does little better. The much larger standard errors of the surprise specification preclude rejection of either extreme hypothesis. These equations are the first ones in this paper to show severe parameter instability. Surprisingly, given the specification and the recent behavior of the government budget deficit, it is the test for stability over half periods, not the test comparing the Reagan period with the rest of the sample, that fails miserably. Tests of the Barro hypothesis over a sample period ending in 1981:2 yield results similar to those in table 7. Finally, we note that the addition of the Barro variables changes a number of the other coefficients. But, to conserve space, we do not detail all these changes.

TAX DISCOUNTING WITHIN A YEAR?

The standard PIH-LCH without bequests can be thought of as the hypothesis that households smooth consumption over their lifetimes. Similarly, the Barro hypothesis can be thought of as the hypothesis that dynasties smooth consumption over periods much longer than a lifetime. Our baseline specification, regressions 3.3 and 3.4, supports one important implication of the standard PIH: that only unanticipated movements in income and wealth lead to changes in consumption. But our results in

44. The government debt variable in the model is real debt per capita, which naturally falls as population and prices rise.

45. Among the other variables held constant is wealth exclusive of government debt, which falls if spending rises. We cannot allow for this reaction, however, because our measure of consumption excludes durables. Hence, the numbers in table 8 slightly overstate the spending stream.

Table 8. Effects on Consumption of a Debt-Financed Temporary Tax Cut, 1981:4–1984:4^a

Dollars

Quarter	No-surprise equation		Surprise equation			
	Estimated effects	Traditional hypothesis	Surprise tax cut		Expected tax cut	
			Estimated effects	Traditional hypothesis	Estimated effects	Traditional hypothesis
1981:4	11.39	11.52	9.88	10.37	2.65	8.79
1982:1	-0.11	2.04	6.52	6.24	1.24	4.35
1982:4	1.67	1.27	2.72	4.74	1.11	2.29
1983:4	2.34	0.92	0.53	3.82	1.00	1.10
1984:4	2.46	0.81	-0.31	3.37	0.97	0.60

Source: Authors' computations based on the equations in table 7.
 a. Real per capita consumption of nondurable goods and services.

table 7 offer at least some evidence against the Barro equivalence hypothesis.

A test of a much weaker smoothing hypothesis—that consumers smooth their spending over the annual April 15 tax settlement date—provides a useful cross check on our previous results. That consumers do at least this much smoothing may seem self-evident, but newspaper reports last winter claimed that consumer spending was low in 1985:1 because of Internal Revenue Service delays in processing tax refunds. We can test the “smoothing over April 15” hypothesis because, while official NIPA data use income tax payments rather than income tax liabilities in defining disposable income, the BEA also publishes a quarterly time series on tax liabilities.⁴⁶

A weak requirement of the PIH ought to be that consumers base their spending plans on income net of tax liabilities, not net of payments. When *TP* is defined as payments and *TL* as liabilities, the concept of labor income relevant to consumers is:

$$YL + \lambda(TP - TL),$$

rather than just *YL*. The assumption is that labor's share in tax liabilities

46. The sample period is one year shorter for these tests because 1984 data on tax liabilities are not yet available. Data sources are as follows: for 1953–75, U.S. Bureau of Economic Analysis, *Survey of Current Business*, vol. 58 (May 1978); for 1976–79, *Survey of Current Business*, vol. 63 (January 1983); for 1980, *Survey of Current Business*, vol. 64 (April 1984); for 1981–83, *Survey of Current Business*, vol. 65 (May 1985).

is the same as labor's share in tax payments. To test whether it is liabilities or payments that really matter, we replace $\ln YL$ wherever it appears by:

$$\ln[YL + \mu_2 \lambda(TP - TL)] \approx \ln YL + \frac{\mu_2 \lambda(TP - TL)}{YL}.$$

The test, however, is not likely to have much power because the variance of $\lambda(TP - TL)$ is small compared to that of YL .

Table 9 summarizes the results. Of the four implied estimates of μ_2 , all but one, which is very poorly determined, are fairly close to unity. Although the formal tests cannot reject either extreme hypothesis, $\mu_2 = 0$ or $\mu_2 = 1$, owing to large standard errors, they prefer the latter. Given that low discriminatory power was to be expected, there appears to be no evidence against short-run smoothing here; so we are inclined to accept the theoretical prediction that liabilities, not payments, influence consumption.

Conclusion: What Have We Learned?

The hypotheses we have tested suggest one main change in our baseline consumption functions: that the income derived from temporary tax changes should be eliminated from disposable labor income. In our original baselines in table 3, we made this adjustment for the 1975 rebate, but not for the other temporary tax episodes. The regressions reported in table 10 make this change.

For several reasons, we tend to think of the surprise regression in table 10 as our best consumption function. First, the "no decomposition" hypothesis that leads to the no-surprise specification is soundly rejected. Second, the surprise regression has smaller unexplained time trends. Third, and related to this, the long-run elasticity of consumption with respect to income, when wealth effects are included by keeping the ratio of wealth to income constant, does not differ significantly from unity in the surprise regression but does in the no-surprise regression. However, the corresponding no-surprise regression is also provided for readers who prefer a more traditional specification that does not rely on an admittedly questionable procedure for dividing variables into anticipated and unanticipated components.

Table 9. Tests of Tax Discounting Within the Year, 1954:1–1983:4^a

Variable, parameter, and test	Regression	
	Surprise	No-surprise
Tax payments minus liabilities income variable ^{b,c}		
Lagged	0.043 (0.2)	–0.004(–0.03)
Anticipated	–0.207(–0.6)	0.155 (1.4)
Unanticipated	0.164 (1.4)	
<i>Parameter</i> ^d		
Estimated μ_2		
Impact	0.84	0.83
Impact and lagged	–2.29	1.37
<i>Test of hypotheses</i> (marginal significance levels) ^e		
$\mu_2 = 0$	0.47	0.31
$\mu_2 = 1$	0.62	0.78
Unit elasticity	0.44	0.14
Real interest rate	0.004	0.076
No decomposition	0.018	...
No decomposition for income and wealth	0.041	...
Surprises only	0.94	...
Half-sample stability	0.23	0.30
Reagan stability	0.57	0.52

Sources: Same as table 2.

a. Dependent variable is the change in log of consumption of nondurable goods and services. For each variable, x_t , the anticipated variable is defined as Ex_t , the unanticipated variable as $x_t - Ex_t$. New variables are added to baseline equations 3.3 and 3.4.

b. The variable is defined as $\lambda[(TP-TL)/YL]$, where λ is labor share, TP is tax payments, TL is tax liabilities, and YL is labor income.

c. These are coefficients of the indicated variables in the regressions.

d. As described in the text, the parameters are estimated as the ratio of coefficients from the regressions.

e. This is the probability, expressed as a decimal, of getting the indicated coefficients if the hypothesis is correct.

The main conclusions of this study are the following.

There do seem to be empirical gains from decomposing income and wealth changes into anticipated and unanticipated components, as the “rational expectations” approach to the consumption function suggests. It seems to be mainly unexpected, not expected, changes in income and wealth that cause consumption to change, just as modern versions of the permanent income hypothesis suggest. In fact, we cannot reject the hypothesis that lagged and anticipated income and wealth are irrelevant to changes in consumption. This may reflect the fact that in our formulations of expectations, much of the time series variance of changes in income and wealth is unexpected.

The rate of interest has an insignificantly negative influence on

Table 10. Augmented Baseline Regressions^a

<i>Independent variable, summary statistic, and test of hypothesis</i>	<i>Regression</i>	
	<i>Surprise 10.1</i>	<i>No-surprise 10.2</i>
Constant	-0.19 (-0.4)	0.38 (1.4)
Time trend ($\times 10^3$)	0.26 (1.0)	0.49 (2.6)
Lagged consumption	-0.126(-1.5)	-0.221(-3.6)
Income		
Lagged	0.141 (0.8)	-0.053(-1.0)
Anticipated	-0.048(-0.3)	0.171 (3.4)
Unanticipated	0.178 (3.3)	
Wealth		
Lagged	0.034 (0.7)	-0.020(-0.6)
Anticipated	0.012 (0.2)	0.060 (2.0)
Unanticipated	0.116 (3.0)	
After-tax nominal interest rate, lagged	-0.274(-1.3)	-0.141(-1.0)
Inflation		
Anticipated	-0.928(-2.7)	-0.259(-1.8)
Unanticipated	-0.127(-0.8)	
Relative price of nondurable goods		
Lagged	...	0.281 (2.8)
Anticipated	-0.135(-2.7)	-0.220(-2.1)
Unanticipated	-0.231(-2.1)	
Relative price of durable goods		
Lagged	...	0.042 (0.5)
Anticipated	0.078 (2.1)	-0.009(-0.1)
Unanticipated	0.022 (0.2)	
Stock of durable goods		
Anticipated	-0.003(-0.1)	-0.001(-0.1)
Unanticipated	0.108 (0.6)	
<i>Summary statistic</i>		
R^2	0.541	0.515
Standard error ($\times 100$)	0.340	0.343
Sum of square residuals ($\times 100$)	0.122	0.129
Degrees of freedom	106	110
Long-run income elasticity	1.08	0.71
<i>Test of hypotheses</i> <i>(marginal significance levels)^b</i>		
Unit elasticity	0.86	0.081
Real interest rate	0.008	0.038
No decomposition	0.015	...
No decomposition for income and wealth	0.081	...
Surprises only	0.81	...
Half-sample stability	0.063	0.14
Reagan stability	0.31	0.34

Sources: Same as table 2.

a. For each variable, x_t , the anticipated variable is defined as Ex_t , the unanticipated variable as $x_t - Ex_t$.

b. This is the probability, expressed as a decimal, of getting the indicated coefficients if the hypothesis is correct.

consumption. What effect it has is on services, rather than on nondurable goods. Surprisingly, it is the nominal, not the real, interest rate that seems to matter. Inflation has a substantial independent influence on consumption. Other things being equal, higher anticipated inflation leads to lower spending.

Relative prices matter. In particular, a rise in the relative price of nondurable goods leads to slower growth in spending on nondurable goods and services in the short run. Relative prices also matter in another important sense: including them in the consumption function changes our assessment of the effects of several critical variables, such as inflation, interest rates, and anticipated income.

Temporary tax changes appear to have had little, if any, effect on consumption, as the PIH suggests. But consumers seem not to have treated the Reagan “temporary tax hike” as a temporary tax.

There is some inconsistency in our equations regarding the Barro neutrality hypothesis. Though it appears that consumers discount the future tax liabilities implied by government debt, it also appears that government bonds are viewed as net wealth. When the income and wealth coefficients are combined, however, the equations clearly imply that a debt-financed tax cut raises consumption substantially.

Simple specifications that relate consumption to income and, perhaps, to one or two other variables, often give misleading results compared with the fuller models discussed here. Several of the preceding conclusions would not be apparent in the simple consumption functions that frequently appear in the literature.

APPENDIX

Time Series on Unanticipated Variables

THE TIME SERIES on unanticipated income, wealth, and inflation used in our original baseline model (regression 3.3) are given in columns 2, 4, and 6 in table A-1. Each series is juxtaposed against the actual change in the corresponding time series, in columns 1, 3, and 5. Changes are in percentages, at quarterly rates.

Table A-1. Income, Wealth, and Inflation, 1954-84

Percent

	(1) <i>Change in log income</i>	(2) <i>Unantici- pated income</i>	(3) <i>Change in log wealth</i>	(4) <i>Unantici- pated wealth</i>	(5) <i>Change in inflation</i>	(6) <i>Unantici- pated inflation</i>
1954:1	-0.866	-0.345	0.650	-0.726	0.509	0.315
1954:2	-0.810	-0.866	1.356	-0.688	-0.287	0.199
1954:3	0.799	0.603	2.156	0.127	-0.907	-0.776
1954:4	1.460	0.452	1.853	-0.001	0.654	0.003
1955:1	0.428	-0.507	0.924	-1.009	0.696	0.520
1955:2	1.941	1.400	1.572	-0.003	-0.484	-0.175
1955:3	1.033	0.259	1.867	0.585	0.303	0.119
1955:4	1.078	0.471	1.858	0.648	-0.546	-0.378
1956:1	0.608	-0.143	-0.238	-0.855	0.384	-0.098
1956:2	0.183	-0.230	0.430	0.705	0.391	0.200
1956:3	0.049	0.110	0.644	0.523	0.018	0.023
1956:4	0.641	0.872	0.096	-0.043	0.066	0.013
1957:1	-0.329	0.049	-1.046	-0.719	0.085	0.134
1957:2	0.029	0.593	0.793	1.500	-0.323	-0.195
1957:3	-0.216	-0.122	-0.073	-0.302	0.305	0.341
1957:4	-0.862	-0.682	-0.833	-0.805	-0.553	-0.316
1958:1	-1.547	-1.267	-1.097	-0.570	0.692	0.521
1958:2	0.360	0.375	1.750	1.241	-0.982	-0.412
1958:3	1.702	0.911	1.322	-0.236	0.130	-0.204
1958:4	1.078	0.306	1.822	0.805	-0.016	-0.249
1959:1	-0.298	-0.369	1.923	1.114	0.486	0.153
1959:2	1.006	1.024	1.720	0.341	-0.136	-0.100
1959:3	-1.198	-1.422	0.621	-0.137	0.273	0.192
1959:4	-0.031	0.063	0.676	0.544	-0.476	-0.291
1960:1	-0.009	-0.512	-0.904	-0.756	0.005	-0.144
1960:2	0.504	0.387	-0.032	0.646	0.207	0.158
1960:3	-0.610	-1.136	-0.350	-0.899	-0.217	-0.177
1960:4	-0.785	-1.446	-0.326	-0.551	0.160	0.249
1961:1	0.518	-0.347	1.379	0.906	-0.524	-0.228
1961:2	1.131	-0.105	2.238	1.114	0.124	-0.042
1961:3	0.741	-0.315	0.808	-0.811	0.350	0.207
1961:4	1.252	0.838	1.657	1.054	-0.298	-0.206
1962:1	0.729	0.152	0.290	-0.500	0.349	0.118
1962:2	0.510	0.171	-1.526	-1.702	-0.043	-0.017
1962:3	-0.246	-0.594	-2.116	-1.518	-0.140	-0.097
1962:4	0.009	-0.429	-0.121	0.418	0.163	0.302
1963:1	0.306	-0.580	1.478	0.199	-0.071	0.277
1963:2	0.596	-0.511	2.083	0.089	-0.196	0.165

Table A-1. Income, Wealth, and Inflation, 1954-84 (Continued)

Percent

	(1) <i>Change in log income</i>	(2) <i>Unanti- cipated income</i>	(3) <i>Change in log wealth</i>	(4) <i>Unanti- cipated wealth</i>	(5) <i>Change in inflation</i>	(6) <i>Unanti- cipated inflation</i>
1963:3	0.503	-0.763	0.649	-1.281	0.248	0.522
1963:4	0.839	-0.449	0.554	-0.511	-0.042	0.396
1964:1	1.652	0.593	1.195	0.330	-0.000	0.362
1964:2	2.637	1.379	1.835	0.579	-0.297	-0.024
1964:3	1.016	-0.393	1.163	-0.078	0.224	0.170
1964:4	0.773	-0.141	1.491	0.676	-0.116	-0.127
1965:1	0.670	0.002	0.790	0.223	0.282	0.057
1965:2	0.714	0.357	0.860	0.590	0.147	0.042
1965:3	2.192	1.238	0.592	0.243	-0.185	-0.116
1965:4	1.506	0.523	1.142	1.183	-0.022	-0.266
1966:1	0.610	-0.324	0.003	-0.324	0.472	0.024
1966:2	0.544	-0.048	-1.398	-1.174	-0.157	-0.287
1966:3	1.079	0.112	-2.565	-1.780	0.047	-0.111
1966:4	1.112	0.390	-1.397	-0.198	-0.124	-0.213
1967:1	1.058	0.197	1.856	2.006	-0.472	-0.528
1967:2	0.541	-0.573	3.321	1.547	0.333	-0.122
1967:3	0.554	-0.347	1.392	-1.022	0.304	0.025
1967:4	0.667	-0.144	0.147	-0.485	0.015	-0.044
1968:1	1.271	0.859	-0.571	-0.294	0.393	0.263
1968:2	1.593	1.167	2.018	2.200	-0.334	-0.222
1968:3	0.265	-0.684	1.905	0.970	-0.001	-0.144
1968:4	0.518	-0.181	1.539	0.731	0.205	-0.003
1969:1	-0.123	-0.880	-0.385	-0.548	-0.136	-0.163
1969:2	0.606	0.058	-1.078	0.065	0.286	0.127
1969:3	1.700	0.826	-1.202	-0.125	-0.119	0.044
1969:4	0.209	-0.589	-1.048	0.276	0.058	0.113
1970:1	0.368	-0.018	-0.893	-0.216	-0.092	-0.007
1970:2	1.846	1.192	-1.825	-1.758	-0.107	-0.109
1970:3	0.489	-0.172	-1.637	-1.211	-0.114	-0.157
1970:4	-0.757	-1.301	1.024	0.635	0.429	0.363
1971:1	1.676	1.125	3.147	0.558	-0.357	-0.017
1971:2	1.170	0.355	2.855	-0.521	0.108	-0.003
1971:3	-0.211	-0.856	-0.107	-2.213	-0.077	-0.067
1971:4	0.520	0.472	-0.073	-0.080	-0.284	-0.383
1972:1	0.699	-0.016	1.584	0.969	0.295	0.049
1972:2	1.021	-0.140	2.815	0.960	-0.283	-0.258
1972:3	0.650	-0.681	1.157	-0.673	0.172	-0.154
1972:4	2.937	1.814	1.882	1.269	-0.005	-0.262

Table A-1. Income, Wealth, and Inflation, 1954-84 (Continued)

Percent

	(1) <i>Change in log income</i>	(2) <i>Unantici- pated income</i>	(3) <i>Change in log wealth</i>	(4) <i>Unantici- pated wealth</i>	(5) <i>Change in inflation</i>	(6) <i>Unantici- pated inflation</i>
1973:1	1.564	0.161	1.494	0.701	0.508	0.023
1973:2	0.521	-0.037	-0.257	-0.505	0.519	0.242
1973:3	0.122	0.096	-0.501	0.431	-0.205	-0.143
1973:4	0.258	0.619	-0.338	1.113	0.307	0.010
1974:1	-2.068	-1.621	-2.003	-1.555	0.887	0.740
1974:2	-0.736	0.497	-1.432	-0.440	-0.354	0.056
1974:3	-0.438	0.365	-2.694	-1.975	-0.118	0.109
1974:4	-1.210	-0.781	-2.618	-1.874	0.044	0.536
1975:1	-0.762	-0.661	1.134	0.827	-1.116	-0.295
1975:2	1.647	1.292	2.969	0.768	-0.183	-0.222
1975:3	1.748	0.420	0.875	-2.070	0.665	0.616
1975:4	0.709	0.133	0.399	-0.884	-0.390	-0.092
1976:1	1.400	1.279	3.179	2.190	-0.544	-0.633
1976:2	0.222	-0.492	3.661	1.268	0.026	-0.411
1976:3	0.173	-0.358	1.761	-0.191	0.403	0.022
1976:4	0.229	-0.262	0.440	-0.419	0.194	0.042
1977:1	0.026	-0.585	-0.759	-1.233	-0.022	-0.121
1977:2	1.238	0.406	-0.079	-0.016	-0.211	-0.287
1977:3	1.485	0.183	0.027	-0.181	0.081	-0.021
1977:4	1.133	0.097	0.252	-0.058	-0.033	-0.083
1978:1	0.271	-0.532	0.321	-0.192	0.194	0.079
1978:2	0.549	-0.180	1.393	0.519	0.665	0.663
1978:3	0.281	-0.270	2.453	0.759	-0.325	0.079
1978:4	0.369	-0.184	1.869	0.245	0.096	0.129
1979:1	0.059	-0.221	0.506	-0.448	0.227	0.235
1979:2	-0.270	-0.584	0.754	0.573	-0.016	0.139
1979:3	0.316	0.230	0.313	0.020	0.041	0.079
1979:4	-0.732	-0.716	-0.064	-0.007	0.128	0.065
1980:1	-0.449	-0.290	-0.019	0.229	0.310	0.326
1980:2	-1.910	-1.598	0.242	0.170	-0.312	0.020
1980:3	0.573	0.945	1.367	0.440	0.015	0.145
1980:4	0.370	0.364	1.584	0.278	-0.108	-0.027
1981:1	-0.162	0.191	0.695	0.828	-0.194	-0.153
1981:2	-0.949	-0.258	0.655	1.031	-0.234	-0.295
1981:3	0.492	1.132	-0.081	0.207	0.056	0.034
1981:4	-0.666	-0.095	-0.356	0.332	-0.194	-0.056
1982:1	-0.804	-0.207	-0.284	-0.358	-0.366	-0.268
1982:2	-0.133	0.373	-0.727	-0.078	-0.360	-0.368

Table A-1. Income, Wealth, and Inflation, 1954–84 (Continued)

Percent

	(1) <i>Change in log income</i>	(2) <i>Unantici- pated income</i>	(3) <i>Change in log wealth</i>	(4) <i>Unantici- pated wealth</i>	(5) <i>Change in inflation</i>	(6) <i>Unantici- pated inflation</i>
1982:3	0.398	0.302	-1.093	-1.003	0.506	0.516
1982:4	0.777	0.113	1.893	0.635	-0.515	-0.004
1983:1	0.302	-0.712	2.764	0.148	-0.422	-0.400
1983:2	0.818	-0.256	1.708	-0.569	0.445	0.290
1983:3	1.352	-0.000	1.614	-0.132	-0.088	0.052
1983:4	1.708	0.470	1.077	-0.165	-0.335	-0.265
1984:1	1.474	0.421	-0.150	-0.747	0.353	0.158
1984:2	0.773	0.069	-0.160	0.035	-0.386	-0.316
1984:3	0.055	-0.483	-0.724	-0.452	0.527	0.296
1984:4	0.529	0.156	0.734	0.558	-0.440	-0.118

Comments and Discussion

Robert E. Hall: Alan Blinder and Angus Deaton make a brave assault on some of the major outstanding issues in aggregate consumption. Unhappily, the data do not always yield up definitive answers to the questions the authors pose. But their efforts convince me that no additional econometric wizardry will get the answers out of U.S. time series data on major categories of consumption.

One of the questions investigated by Blinder and Deaton has received a great deal of attention already, most notably in work by Marjorie Flavin.¹ Can the first difference of consumption be predicted, or is it purely a measure of the consumer's reaction to new, inherently unpredictable information? Blinder and Deaton replicate Flavin's finding that the change in consumption is predictable from past income. In their equation 2.1, from table 2, lagged income and the forecast of current income based on lagged information are both significant, contrary to the implications of a simple rational expectations model. However, in a more elaborate rational expectations model, expressed in equation 3.3, from table 3, with interest rates and relative prices, the predictive power of lagged income and wealth disappears. Flavin's rejection of rational expectations may be an artifact of her neglect of predictors that are consistent with rational expectations. However, Blinder and Deaton do not pursue this finding any further.

Blinder and Deaton are concerned with another hypothesis they call "no decomposition." Under it, the influence of the surprise component of, say, income is the same as the influence of the predictable component.

1. Marjorie A. Flavin, "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy*, vol. 89 (October 1981), pp. 974-1009.

In that case, income itself is the proper variable, and the decomposition into surprise and predicted components is irrelevant. Blinder and Deaton are extremely unclear as to why we should be interested in the hypothesis. They justify testing the hypothesis on the grounds that “the PIH suggests so strongly that anticipated and unanticipated income and wealth should get different coefficients.” However, they acknowledge in a footnote that “under some circumstances, the pure PIH is fully consistent with our no-surprise specification.” In other words, sometimes a rational expectations model satisfies the no decomposition hypothesis, and sometimes it does not, depending on issues that have nothing to do with rational expectations. The paper is devoid of any analysis supporting the proposition that rational expectations, or any other interesting proposition, is related to the no decomposition hypothesis. The work that I have done suggests that no interest attaches to the no decomposition hypothesis. I think the paper should stick to testing hypotheses that have been carefully derived from theory, such as the surprises-only hypothesis.

The role of interest rates in the consumption function has received a good deal of attention recently, both within and outside the rational expectations framework. Blinder and Deaton find slightly negative coefficients, but cannot reject the hypothesis of no effect. The finding is confused somewhat by the strong negative effect of expected inflation on consumption. Plainly, expected inflation is not playing the role of making only the real rate of interest affect consumption. Blinder and Deaton attribute the negative role of inflation to consumer confusion regarding real and nominal changes. I think they should also consider the proposition that inflation is a specific symptom of unfavorable events in the U.S. economy. Over their sample period the two biggest bursts of inflation occurred following oil price shocks. Their results say that consumers react more strongly to an oil shock than they do to other events with the same impact on real income and real wealth.

When Blinder and Deaton turn their attention to the important issue of the effect of temporary tax measures on consumption, they are again unable to coax a definite answer out of the data. A temporary tax cut has virtually no immediate effect on consumption, as predicted by permanent income-rational expectations theory, but the effect in the long run is for consumption to fall, which is inconsistent with any theory. Blinder and Deaton also make an ingenious analysis of the phased-in tax cuts adopted

in 1981. Their results suggest that consumers waited until the cuts went into effect to start spending them. This finding is consistent either with a simple view that consumers respond, irrationally, only to their current disposable income, or with the view that consumers are rationally skeptical about the likelihood of promised future cuts.

The last major topic is the Barro equivalence hypothesis, which states that consumers hold back during periods of budget deficits because they know they will have to pay extra taxes later to support the resulting government debt. Blinder and Deaton set up the test in a way that looks separately at an adjustment to income (subtracting the deficit from income) and an adjustment to wealth (subtracting government debt from total wealth). They find that consumers do seem to use the income adjustment; deficits do reduce consumption via this channel. This finding is puzzling in light of table 1, which shows that private saving has not risen substantially during the period of huge deficits starting in 1982. Part of the answer comes from their other finding that consumers do not ignore government debt in calculating their wealth, as they should according to Barro. Taken together, the net effect of an increase in the deficit working on consumption through both the labor income and wealth (nonlabor income) channels is to increase consumption.

R. Glenn Hubbard: Reverence for “psychological laws” notwithstanding, the regularity of the relationship between consumption and income must surely strike us as a conundrum. Exploration of this relationship—that is, of the time series behavior of aggregate consumption—has been a focus of macroeconomic research for decades. In recent years, theoretical and empirical efforts have pursued two related avenues of inquiry. One is the extent to which “consumption functions” are consistent with the predictions of the life-cycle model or the permanent income hypothesis. The other is the “effectiveness” of fiscal policy changes in raising or lowering consumption.

Blinder and Deaton have given us a careful and innovative reexamination of both issues. The paper presents a systematic evaluation of recent empirical propositions regarding consumption models or, more correctly, models of consumer spending on nondurables and services. By addressing a large set of variables potentially affecting consumption, their study offers a rich comparison with previous studies. Two issues are particularly important: the notion from some previous studies that

only unanticipated income changes matter, and the hypothesis associated with Barro that individuals pierce the veil of government finance. My focus is on those two issues and, more generally, on the relevance of various modeling strategies for policy analysis.

Blinder and Deaton begin with a flexible functional form for estimation and analysis in equation 1. The choice is a convenient one, as the model nests many alternative hypotheses, and as is pointed out, the steady-state value of the average propensity to consume can be solved for to test whether the results are sensible. The paper's reference to an error-correction model raises the questions of what factors are motivating the adjustment process and whether such a process of adjustment might also be influenced by transitory unemployment or liquidity constraints. For example, equation 1 is consistent with an error-correction model in which agents adjust toward a desired consumption-income ratio, as opposed to a desired consumption level. As an empirical matter, based on the estimates in table 2, the speed of adjustment of the consumption-income ratio toward its desired level is quite slow.

The results in table 2 are an informative expansion of "traditional" consumption functions. The addition of household net worth variables improves the fit of the regression, as one would expect. What concept of "net worth" to use here is a tricky question. For example, should one use only financial net worth? More important, what should be done about household claims to social security and private pension benefits? The former is unfunded, and the latter is not, but both clearly dominate household wealth. The issue is important in trying to disentangle effects of "news about future wealth" from those of changes in liquidity. A (log) wealth-income ratio variable could reflect either a scale parameter for $\ln(C/Y)$ or "precautionary" saving against future events leading to liquidity constraints on consumption. Its expected effect is unclear.

A time trend is also included when the household wealth variables are introduced; I am not sure of its interpretation in equations 2.3 and 2.4. It is true that C here is not total consumer expenditures and that component parts might exhibit "trends" if income and wealth were the only explanatory variables. However, the inclusion of relative prices of, say, durable and nondurable goods should mitigate that problem. When this adjustment is made, in table 3, the time trend is still generally statistically significant. One can imagine other "trend" influences on consumption relative to income (for example, the increasing generosity of social insurance programs) for which the time trend may be a proxy.

Perhaps most interesting in table 2 are the results for models including labor income, the concept suggested by the life-cycle model. This first set of results illustrates that only actual variables, not their division into anticipated and unanticipated components, matter. Even in this simple case, however, one caveat is that consumers' expectations about income or net worth positions may be substantially more informed than those of the econometrician, so that part of the difference between "actual" and "expected" variables is not really unexpected. The larger is that source of "error" relative to the "true" forecast error, the more likely are the estimates to give the result that "only actual variables matter."

The addition of relative price effects, both intertemporal and across categories of consumption, in table 3 brings still more novel, and sometimes puzzling, findings. Nominal after-tax interest rates exert the predicted negative effect on consumption, though the coefficient estimate is statistically insignificant. That the expected inflation variable does not have the positive sign required by a "real interest rate" explanation could reflect a relative shift to spending on durables in times of inflation. Based on the results in equations 3.3 and 3.4, the "price confusion effect" noted by the authors is suspect, given the importance of anticipated inflation relative to unanticipated inflation. It is more likely that effects of recurrent episodes of inflation on consumption reflect the importance of the structural demand and supply variables underlying inflationary periods.

Inclusion of relative prices of durable and nondurable goods gives the expected effects and improves the fit of the model. In addition, both the hypotheses that "only actual variables matter" and "only surprises matter" can be rejected. The finding that nominal interest sensitivity of spending on services exceeds that of spending on nondurable goods is indeed unusual. That services are relatively more complementary with durable goods seems unlikely. The results on the interest sensitivity of consumer spending do not shed much additional light on the debate over the sign and magnitude of the interest elasticity of saving. In addition to the problems of carrying out such an analysis with aggregate time series data, theory tells us that there can be no one "interest elasticity of saving"; results are specific to the individual thought experiment.

Perhaps most important for the application of the aggregate consumption function to policy analysis is the issue of its ability to assess whether households internalize the saving decisions made for them by the

corporate and government sectors. An obvious starting point for research is to test for differences in consumer responses to temporary and permanent tax changes. The Blinder-Deaton paper follows up Blinder's important earlier study and presents, among other things, an early look at the effects of the 1981 Economic Recovery Tax Act.¹ While the results are not precise, it appears that scheduled tax reductions during the Reagan period were ignored by individuals. That is, individuals consumed according to actual after-tax incomes rather than the permanent after-tax incomes they would have once the tax reductions were fully effective. The estimates given in this paper for temporary tax changes in the pre-Reagan period differ from those in Blinder's previous study; here, the short-run effects of temporary taxes correspond more closely to those predicted by the permanent income hypothesis.

The authors do not consider whether individuals "pierce the corporate veil." Nor do they question whether individuals pierce their own veil, that is, the extent to which household net worth held in private pensions, or, for that matter, social security claims, acts like nonpension wealth in its effect on consumption. In other words, is household net worth held in private pensions substitutable for nonpension net worth?

Most significant for the policy debate over Ricardian equivalence and the government debt, Blinder and Deaton examine whether households distinguish between debt and taxes in government finance. The literature gives many reasons why the Barro hypothesis might not be borne out by the data: unlike the stylized model of Ricardian equivalence, tax policy is not carried out on a lump-sum basis; bequests may be at a corner solution for many if not most consumers; in the presence of capital-market imperfections, binding liquidity constraints on many individuals may restrict consumption movements. The last explanation seems particularly relevant here, given some of the other evidence in the paper. The authors do find evidence for discounting of future tax liabilities, though it is not supported by the wealth coefficients in table 7. The estimates are not precise enough to draw specific conclusions. Despite their rejection of the strict requirements of the Barro model, Blinder and Deaton find evidence for short-term smoothing, evidence required for even a pragmatic belief in life-cycle or permanent income models.

1. Alan S. Blinder, "Temporary Income Taxes and Consumer Spending," *Journal of Political Economy*, vol. 89 (February 1981), pp. 26-53.

A more direct way to contrast the Barro model and the life-cycle model with time series data is to test whether consumption varies with the age distribution of resources. Unlike the life-cycle model, the Barro model predicts that consumption is a function only of collective resources, not of the age distribution of resources. Data on the age distribution of property and labor income are available, and one can test whether “income share” variables matter in the consumption equation. As tax policies frequently redistribute income and wealth along these lines, this is likely to be a fertile area for applied research. In addition, such tests are likely to be important in the continuing debate over the influence of intergenerational debt policy on consumption.

Where do we go from here? By now, two puzzles figure prominently in recent applied research on consumption. First, given the professional consensus on the modeling paradigm for the life-cycle and permanent income hypotheses, why do versions of “traditional” consumption functions seem to fit the data well? Second, why is it that estimates of aggregate saving generated from theoretically believable life-cycle consumption simulation models fall short of “realistic” values?²

I suspect that these puzzles, though revealed in different research agendas, are related. Two modifications of existing work are suggested. First, we need to consider more explicitly how restrictions on private trades—specifically, borrowing restrictions, or liquidity constraints—affect consumer spending, particularly in extensions to models of spending on durables. For example, life-cycle microsimulation models can be used to examine the sensitivity of aggregate consumption and the capital stock to liquidity constraints, helping us to see whether likely effects are sizable enough to explain the “anomalies” in aggregate consumption functions. Second, households most certainly engage in precautionary saving—against uncertainty over earnings, health, length of life, and so forth. Dynamic life-cycle simulation models may be able to shed light here as well. To the extent that precautionary motives in response to individual-specific uncertainty are important, it may be difficult to rationalize results of aggregate consumption functions with aggregated

2. See, for example, Laurence J. Kotlikoff and Lawrence H. Summers, “The Role of Intergenerational Transfers in Aggregate Capital Accumulation,” *Journal of Political Economy*, vol. 89 (August 1981), pp. 706–32; and R. Glenn Hubbard and Kenneth L. Judd, “Social Security and Individual Welfare: Precautionary Saving, Liquidity Constraints, and the Payroll Tax” (Northwestern University, 1985).

results of simulation models of individual behavior. These two qualifications of the basic model—restrictions on borrowing and precautionary saving—are certainly related.

The Blinder-Deaton paper is an important integration of various lines of inquiry in aggregate time series studies of consumption. New research on issues involved in linking solutions to individual optimizing problems to aggregate data (in particular, the modeling of liquidity constraints) will be an important step in extending their work.

General Discussion

Franco Modigliani questioned Alan Blinder and Angus Deaton's exclusion of durable goods services from their consumption measure; he argued that they ought to have imputed rental values to the services of durable goods and included those imputed values along with purchases of nondurables and services in total consumption. Blinder defended the consumption measure chosen, noting that durable goods services may not be perfect substitutes for other forms of consumption, as would have to be assumed to simply add them to the total, and that the dynamics of durable goods purchases are likely to differ considerably from the dynamics of other consumption purchases.

There was general discussion of whether significant "anticipated" right-hand side variables are inconsistent with a model in which consumption depends only on permanent income, itself a random walk. Christopher Sims observed that one explanation sometimes given for the significance of anticipated variables is the presence of "transitory consumption." But either transitory consumption reflects information available to consumers but not available to economists modeling their behavior—in which case changes in it should not be serially correlated so that its presence cannot explain significant coefficients on anticipated variables—or it reflects an inadequacy of the theory—in which case it could be serially correlated, but no longer satisfies the identifying assumptions of the random walk model. In this last case, simultaneous equation methods are necessary to test the theory.

Sims also urged that more careful attention be given to time aggregation issues, which are critical if one is testing dynamic propositions related to the importance of innovations. The predictions emerging from

rational expectations theory can look quite different when they are based on data aggregated over two or three time periods, as when monthly data are aggregated into quarters, than they do when they are based on data for single periods. Stephen Goldfeld noted that shifts in the age and income distribution of the population could introduce errors into the aggregate consumption functions Blinder and Deaton estimate.

Stanley Fischer commented that the categorization of income changes as expected or unexpected depends critically upon the underlying model of expectations. Insofar as these distinctions are crucial to the analysis, he urged that more attention be paid to the vector autoregressions that produce the income change predictions. He also conjectured that in the future, economists will place less reliance on purely econometric models of expectation formation and rely increasingly on data from surveys in which individuals are asked directly about their expectations. Goldfeld noted a conceptual problem that arises whenever aggregate analyses resting upon an assessment of expectations are performed: there will always be some dispersion in expectations within the population, and there is no fully satisfactory rule for deciding which expectations to use.

The question of whether Blinder and Deaton had adequately captured expected and unexpected changes also came up in connection with their test of the intertemporal substitution argument. They reached the puzzling—if by now familiar—conclusion that it is nominal rather than real interest rates that seem to affect consumption. Sims put forward his view that the result reflects errors in the inflation prediction that underlies the distinction between real and nominal interest rates. Standard techniques may lead to overly variable inflation forecasts; better forecasts can be obtained using techniques that shrink the variation in the inflation forecasts towards zero. If the inflation forecasts do not vary so much, then surprises in the nominal rate of interest turn out to be largely surprises in the real rate.

Sims also questioned Blinder and Deaton's treatment of the Barro hypothesis that increases in government debt are seen by consumers as equivalent to increases in their own indebtedness. Blinder and Deaton's test is based on a single equation that includes government debt as an explanatory variable; however, if there is a systematic connection between large deficits and other events that change people's expectations about the likely course of future income, results based on this sort of single equation framework could be quite misleading. Maurice Obstfeld

argued that responses to increases in the government deficit will depend upon whether consumers expect any accompanying increase in government spending to be temporary or permanent. An increase in government spending that is expected to be temporary will have a much smaller effect on consumption than will an equal increase that is expected to be permanent.

George von Furstenberg questioned Blinder and Deaton's characterization of the Reagan tax package as a permanent tax cut combined with a temporary tax increase. In making this characterization, Blinder and Deaton assume a particular sort of forward-looking behavior by consumers estimating their future tax liability; a more sophisticated assumption would be that consumers estimating their future tax liability consider not only provisions already written into the tax code but also likely future changes in the code. In von Furstenberg's view, the large current federal deficit should lead consumers to expect future federal tax increases, even though they have not yet been written into law. This could explain why saving rates did not fall following passage of the 1981 tax act.