

How Important Is the Stock Market Effect on Consumption?

Sydney Ludvigson and Charles Steindel

The second half of the 1990s has seen substantial changes in the wealth of American households, primarily owing to movements in the stock market. From mid-1994 to mid-1997, the aggregate value of household sector equity holdings (including those owned by nonprofits, mutual funds, and pensions and other fiduciaries) roughly doubled, for a dollar gain of about \$5.2 trillion.¹ Since then, stock market values on balance have continued to rise, but there have been massive fluctuations within a wide band; the dollar value of movements within the band—from the low in October 1997 to the recent highs—has been greater than \$3.0 trillion.²

These enormous swings in wealth no doubt have major implications for consumer spending. For this reason, the ability to measure the implications of the swings—that

is, to determine their “wealth effect” on consumer resources—has grown in importance with the changing economic environment. In this article, we examine the wealth effect of stock market changes on consumption. Other things equal, an increase in the stock market makes people wealthier. In general, the wealthier people are, the more they spend. Is it possible, then, to quantify these simple truisms and come up with plausible estimates of the extent to which aggregate consumer spending in the 1990s has been supported by increased stock market wealth? Furthermore, how much would a market correction negatively affect future spending?

Our answers to these questions are a bit limited. We find, as expected, a positive connection between aggregate wealth changes and aggregate spending. Spending growth in recent years has surely been augmented by market gains, but the effect is found to be rather unstable and hard to pin down. The contemporaneous response of consumption growth to an unexpected change in wealth is uncertain and the response appears very short-lived. Therefore, we conclude

Sydney Ludvigson is an economist and Charles Steindel a senior vice president at the Federal Reserve Bank of New York.

that forecasts of future consumption growth are not typically improved by taking changes in existing wealth into account.

In the past, uncertainty about both the long-run (or trend) effect of wealth on consumption and the contemporaneous effect was of modest importance. However, in the current economy—where aggregate wealth fluctuations can be very large relative to household income, spending, and GDP—we find that the uncertainty about the size of the wealth effect also adds a considerable amount of uncertainty to one’s ability to understand trends in consumer spending, over and above the difficulty of understanding the forces behind market movements.

In the next section, we briefly review changes in household sector spending, saving, and wealth, and highlight the central importance of stock market fluctuations to cyclical movements in the household balance sheet. We then turn to econometric analysis to measure the effect of a change in wealth on consumer spending. We find that a traditional specification of the consumption function gives a fairly erratic estimate of the wealth effect and may even suggest that the effect was rather small in recent years. By refining the specification and estimation of the consumption equation to reflect more rigorously current econometric concerns, we narrow the estimate somewhat, but are still left with some instability in our result. Using a more up-to-date methodology, we first establish that consumption and wealth, along with labor income, share a common trend. When asset values or labor income rises, consumption tends to rise as well, and we assess the magnitude of this boost to consumption by estimating the parameters of the shared trend—the marginal propensities to consume out of wealth and labor income. Our results suggest that these propensities are somewhat unstable over the postwar period. Nevertheless, we conclude that a dollar increase in wealth likely leads to a three-to-four-cent increase in consumption in today’s economy, consistent with widely held beliefs about the long-run impact of wealth on consumption.

Finally, we analyze the short-run effects of wealth on consumption by investigating the dynamic response of consumption growth to a change in wealth and by testing the predictive power of wealth for changes in consumer spending. We find that changes in wealth are not corre-

lated with the next quarter’s consumption growth and do not help predict the growth in out-of-sample forecasts. The reason for this is not that wealth has no impact on consumption; rather, the response of consumption growth to an unanticipated change in wealth is largely contemporaneous. Controlling for lagged consumption, changes in the growth rate of wealth provide little additional information about the future path of consumption growth.

THE BASICS OF HOUSEHOLD WEALTH ACCUMULATION AND SAVING

In the aggregate, household wealth accumulation reflects two factors: saving from current income and changes in the valuation of previously owned wealth. The second factor completely dominates changes in aggregate wealth in the short and intermediate terms. In turn, changes in the valuation of existing assets are dominated by fluctuations in the stock market. These points are illustrated in Chart 1. The top panel shows, since fourth-quarter 1952, the cumulated

*In the aggregate, household wealth
accumulation reflects two factors: saving
from current income and changes in the
valuation of previously owned wealth.*

value of increases in household wealth and the cumulated value of household capital gains on the stock market (capital gains are measured as the increase in the value of holdings less cumulated purchases of stock; all series are measured in chain-weighted 1992 dollars). The similarity of the two lines over short time periods is striking. The bottom panel plots the correlation between the changes in the two series over intervals from one to forty quarters, and again shows the overwhelming importance of gains and losses in the stock market in explaining movements in aggregate wealth at anything up to the longest frequencies.

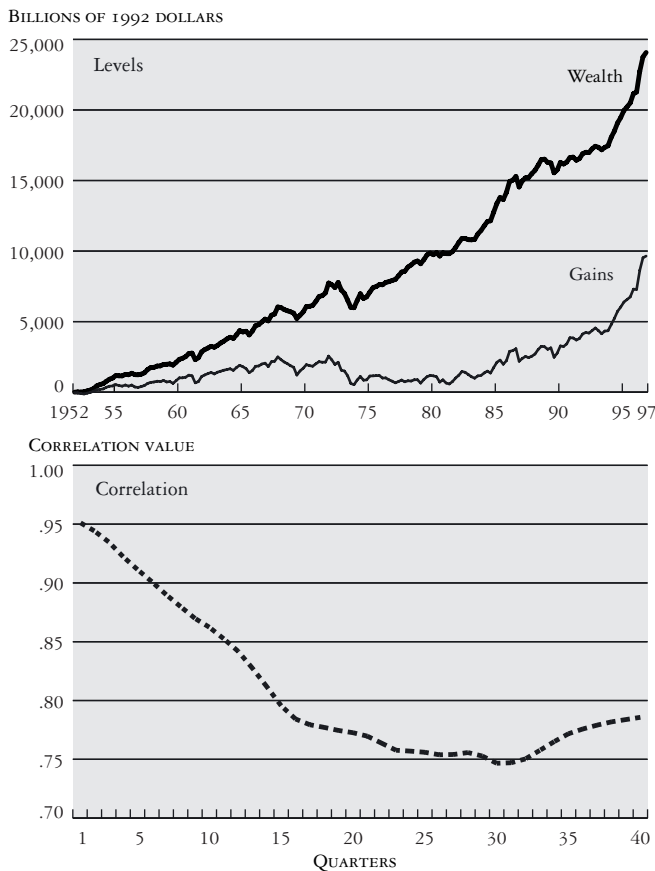
It is clear, then, that in the short run, changes in the pace of wealth accumulation owe little to changes in saving (and other things equal, changes in spending).

However, we are concerned with the opposite issue: the linkage from changes in wealth accumulation to changes in saving and spending.

One way to look at the possible influence of wealth accumulation on saving is shown in Chart 2, which plots the ratio of wealth to disposable income against the personal saving rate. Over the last few years, the wealth-to-disposable-income ratio has increased markedly while the personal saving rate has plunged. The argument for a strong wealth effect is that this increase in the ratio of wealth to disposable income, primarily because of the rise in the stock market, has boosted consumer spending and has reduced saving (both relative to income).

Chart 1

GROWTH IN WEALTH AND CUMULATED STOCK MARKET GAINS




Sources: Board of Governors of the Federal Reserve System; authors' calculations.

Notes: The top panel shows the difference between wealth and the fourth-quarter 1952 level of wealth plotted against the cumulated gains at that point in time. The bottom panel shows the correlation between differences in wealth and differences in cumulated gains over "N" quarters.

Chart 2

WEALTH-TO-DISPOSABLE-INCOME RATIO AND PERSONAL SAVING RATE



Sources: U.S. Department of Commerce, Bureau of Economic Analysis; Board of Governors of the Federal Reserve System; authors' calculations.

However, a simple observation of Chart 2 is not sufficient to establish a well-defined and measured wealth effect. At the most obvious level, the chart shows periods when saving rate moves seem to parallel moves in the wealth-to-disposable-income ratio—for instance, both were increasing in the years around 1980. The seemingly strong negative connection in recent years may be a coincidence. It is helpful to recall that saving is the difference between income and spending. If we are interested in the link between wealth and consumption, it makes more sense to look at consumption directly. Accordingly, we will now turn to a statistical examination of the wealth-spending link.

THE STOCK MARKET AND THE CONSUMER: GENERAL CONSIDERATIONS AND PRELIMINARY EVIDENCE

Traditionally, the wealth effect has been measured by estimating aggregate time-series regressions of the form

$$(1) \quad C_t = a + bW_t + cYP_t + e_t,$$

where C is consumer spending during a period; YP is a measure of permanent income (usually a distributed lag on realized after-tax income); W is consumer net worth, as measured at the beginning of the period; and e_t is an error term capturing other factors that influence consumption.

Derivations of such equations from the underlying theory of consumer behavior may be found in Modigliani

and Tarantelli (1975), Modigliani and Steindel (1977), and Steindel (1977, 1981). The estimated coefficient, b , on wealth, is described as the marginal propensity to consume out of wealth and is interpreted as the increase in consumer spending associated with an increase in wealth. A widespread empirical practice is to separate wealth into different categories, with stock market wealth usually being one

A common assumption is that . . . roughly five cents of each dollar of an increase in wealth is spent soon after it is earned. While this amount seems small, when we are looking at trillion-dollar gains in wealth from the stock market, a five-cent increase in spending per dollar of gain adds up to real money.

of them. A coefficient on stock market wealth different from other types is merely viewed as an artifact of heterogeneity of consumers; stock market wealth owners may be systematically older or younger than other wealth owners or have other characteristics that lead to a different aggregate propensity to consume out of this form of wealth. A common assumption is that b is on the order of .05 or perhaps a bit smaller; in other words, roughly five cents of each dollar of an increase in wealth is spent soon after it is earned. While this amount seems small, when we are looking at trillion-dollar gains in wealth from the stock market, a five-cent increase in spending per dollar of gain adds up to real money.

The perspective of modern dynamic economics is to be quite dubious about the value of estimations such as equation 1, especially using aggregate time-series data. There are questions about the appropriate estimation technique, given the possible presence of aggregation and simultaneity bias, and the use of largely untested simplifying assumptions to derive the estimating equation from the theory. Furthermore, because traditional specifications

and estimation techniques basically assume that consumers are in a steady state, they do not explicitly take into account the adjustment of consumer behavior to new conditions. Formally taking into account the adjustment process to a new equilibrium implies very different ways to specify and estimate the relationship between changes in wealth and changes in consumption. This issue has been addressed in the literature at least going back to Hall (1978).

Despite the valid criticisms of formulations such as equation 1, we establish an initial reference point by estimating this type of model. Equations of this sort have been very influential in the literature on economic policy (see, for instance, Modigliani [1971]) and continue to be common in forecasting exercises.³ Table 1 shows estimates of this traditional type of model. The regressions relate consumer spending to disposable personal income and wealth, with wealth split into two components: stock market holdings and other. Four lags of each of the right-hand-side variables are included in order to capture the adjustment process of consumer spending to changes in fundamentals. Details about the data are provided in Appendix A. The estimation of the model includes a correction for first-order autocorrelation in the error process.

Column 1 shows the coefficients for the equation estimated over the 1953-97 period. The estimates include the sum of the lag coefficients on each of the right-hand-side variables along with the standard errors. These results are more or less consistent with traditional views of consumer behavior: the sum of the lag coefficients on income is roughly .7; the sum of the coefficients on stock market wealth is .04 and, on other forms of wealth, about the same. Each sum is more than twice as great as its computed standard error, which is normally interpreted as meaning that the sum is statistically greater than zero. The estimated coefficient of serial correlation, while substantial, appears to be less than one, suggesting that the model is a valid statistical construct.

The superficial view would be that the equation in column 1 supports traditional opinions of the stock market's impact on consumption. However, the estimated stock market effect appears to be rather sensitive to the period of estimation. Reestimating the equation over three different

periods (columns 2, 3, and 4 of Table 1) suggests that the stock market effect was larger in the late 1970s and early 1980s than either before or after.

Admittedly, columns 2-4 work hard to show this instability. If we divide the sample into three fourteen-year periods (columns 5-7) rather than picking 1975 and 1985 as the break points, the coefficient estimates look more stable, though their standard errors vary. However, Chart 3 reinforces the view of a shifting model. It shows the estimated sum of the lag coefficients, along with one-standard-deviation error bands, of the wealth and income terms from regressions of the form in Table 1 estimated over ten-year periods. In particular, the remarkable thing about the middle panel is not so much the observation that such a parameter changes over time, but that the change from year to year in the estimated effect looks rather large—ten-year regressions estimated ending in two consecutive years will have 80 percent of their observations in common.⁴ The chart also shows that the point estimate of the sum of the lag coefficients on the stock market for the most recent ten-year period is near zero. If all pre-1988 data were destroyed, we would be hard pressed to conclude that there is a link between the stock market and consumer spending, based on this model and estimation technique.

It is clear that the estimated marginal propensity to consume from stock market wealth is not particularly

stable. Of course, it is no great surprise to find uncertainty of this type about a behavioral parameter. The likelihood ratio test statistics reported in Table 1 suggest that we can reject the null hypothesis of a stable structure over the three subsamples in the two parts of the table. In principle, we might try to determine more precisely the break points in the structure of the regression. However, if there is a violation of any of the classical assumptions needed to apply such tests for an equation estimated by ordinary least squares (OLS)—possibilities we discuss further in the next section—the tests of the stability of equation 1 will also be invalid.

Setting aside these concerns, we find that for the purpose of policy analysis, the conventional consumption function estimates produce two important but rather conflicting results. With some trivial exceptions, we consistently come up with estimates of the stock market wealth effect (and the non-stock-market wealth effect) in the range of small positive values to .1—certainly in line with traditional views. Nonetheless, awareness that this propensity can vary in this range makes the wealth effect a very shaky reed to lean on when the aggregate value of the stock market has shown that it can fluctuate by more than \$3 trillion in brief amounts of time. Applying a range of uncertainty about the size of the marginal propensity of only .02 (generally equal to a two-standard-deviation error band for most of our estimates) to such a swing in wealth

Table 1
ORDINARY LEAST SQUARES ESTIMATION OF TRADITIONAL LIFE-CYCLE MODEL

$$\text{Model: } C_t = \sum_{i=0}^3 \delta_i Y_{t-i} + \sum_{i=0}^3 \xi_i SW_{t-i} + \sum_{i=0}^3 \mu_i NW_{t-i} + \epsilon_t$$

Independent Variable	Estimation Period						
	1 1953:1-1997:4	2 1953:1-1975:4	3 1976:1-1985:4	4 1986:1-1997:4	5 1953:1-1967:4	6 1968:1-1982:4	7 1983:1-1997:4
Income (Y)	0.731 (0.067)	0.711 (0.059)	0.568 (0.195)	1.015 (0.077)	0.684 (0.091)	0.832 (0.141)	0.822 (0.074)
Stock market wealth (SW)	0.040 (0.009)	0.026 (0.010)	0.106 (0.041)	0.021 (0.011)	0.030 (0.018)	0.023 (0.019)	0.042 (0.010)
Non-stock-market wealth (NW)	0.038 (0.017)	0.043 (0.015)	0.069 (0.048)	-0.027 (0.017)	0.049 (0.020)	0.012 (0.036)	0.016 (0.018)
Serial correlation coefficient	0.937 (0.030)	0.781 (0.090)	0.937 (0.069)	0.755 (0.097)	0.800 (0.094)	0.886 (0.069)	0.809 (0.091)
Standard error of regression	70.7	59.8	86.7	65.7	41.4	84.7	76.2
Sum of squared residuals of regression	830835	279012	202961	150994	78739	336836	272807
Likelihood ratio test:							
Statistic			48.690			33.668	
Probability			0.0045			0.1436	


Source: Authors' calculations.

Notes: All data are real per capita. Standard errors are in parentheses.

adds \$60 billion (about $\frac{3}{4}$ of 1 percentage point of aggregate GDP) to our uncertainty about the basic forces affecting consumer spending. Table 1 suggests that the range of uncertainty about the wealth propensity should also take into account the different point estimates, which make the range greater than .02. As an extreme example of our

Chart 3

MARGINAL PROPENSITY TO CONSUME



Source: Authors' calculations.

Notes: The panels depict rolling regressions over ten-year samples. The years represent the starting date of each regression. The dashed lines indicate one-standard-deviation error bands.

Table 2
ESTIMATED IMPACT OF 1994-97 STOCK MARKET RISE
ON 1997 CONSUMER SPENDING

Propensity to Consume	Dollar Impact of Wealth Increase on Real Spending	Percentage of 1997 Consumer Spending
0.040 (1953:1-1997:4)	\$166 billion	3.4
0.106 (1976:1-1985:4)	\$439 billion	8.9
0.021 (1986:1-1997:4)	\$87 billion	1.8

Sources: U.S. Department of Commerce, Bureau of Economic Analysis; authors' calculations.

Note: The increase in real household sector stock market holdings, measured from second-quarter 1994 to second-quarter 1997, is \$4,141 billion.

uncertainty about the recent scope of the wealth effect, Table 2 presents a range of estimates of the effect of the 1994-97 stock market rise on the 1997 level of consumer spending. These estimates are taken by applying the propensity to consume from stock market wealth determined from columns 1, 3, and 4 of Table 1 to the rise in the aggregate real value of household sector stock market wealth in this period. The estimated range of the effect of the 1994-97 market rise on 1997 spending spans more than 350 billion chain-weighted 1992 dollars. Alternatively, we can argue that the 1994-97 market increase boosted 1997 spending somewhere between $1\frac{3}{4}$ and 9 percent. Even the smallest effect can account for the 1.5-percentage-point drop in the personal saving rate over that period. However, the range of the estimates is clearly very disquieting. We now turn to more modern statistical techniques to obtain a more precise handle on the wealth effect.

THE WEALTH EFFECT ON CONSUMPTION:
UPDATED STATISTICAL APPROACHES

This section employs updated empirical techniques to investigate the relationship between consumption and wealth. We begin by estimating the marginal propensity to consume out of wealth with more modern econometric procedures. With estimates of the marginal propensity to consume out of wealth in hand, we move on to analyze the response, over time, of consumption growth to a wealth shock, and to test whether accounting for movements in

wealth is likely to improve our forecasts of consumption growth one or more quarters ahead.

LONG-RUN RELATIONSHIPS: THE MARGINAL PROPENSITY TO CONSUME OUT OF WEALTH

The empirical procedure above provides a descriptive summary of the relationship between aggregate consumption and wealth. Studying those results is useful because it furnishes a basis for comparison with earlier work in the traditional life-cycle consumption literature. That empirical

Much recent theoretical research on the consumer has focused on the behavior of a representative individual who is forward looking but faces a risky stream of labor income. Among the most prominent of these paradigms is the permanent income hypothesis.

methodology is still widely used today. Nevertheless, econometric theory points to a number of potential pitfalls with the traditional approach to estimating the effect of wealth on consumption.

One potential pitfall concerns the failure to account for the time-series properties of C , W , and Y . At the least, each of these variables likely contains a trend component that is random and therefore not known in advance (a stochastic trend). The conventional analysis performed above does not take into account the econometric implications of this type of nonstationarity. A second problem pertains to the correlation between consumption and current wealth. We seek to identify the effect of an increase in wealth on consumption. Yet the econometric techniques employed above ignore the possibility that the estimated consumption-wealth correlation reflects, at least partially, the effect of an increase in consumption on wealth.⁵ We refer to this “reverse causality” as endogeneity bias. Failure to address either problem could skew statistical

inference and lead to inconsistent estimates of how much an increase in wealth influences consumption. We now present an alternate empirical approach and discuss how it can address both difficulties.

We begin by laying some theoretical groundwork. Our purpose is solely to provide intuition and motivation for the statistical analysis that follows; the empirical approach we take is not conditional on any particular theory of consumption and will be robust to a variety of departures from the framework presented next. We discuss this further below.

Much recent theoretical research on the consumer has focused on the behavior of a representative individual who is forward looking but faces a risky stream of labor income. Among the most prominent of these paradigms is the *permanent income hypothesis*. According to this theory, consumption of nondurable goods and services is chosen to match permanent income, defined as the annuity value of human and nonhuman wealth. The model implies that consumption responds to any unpredictable change in permanent income, but very little to transitory fluctuations in income. Additionally, there are no lags in the adjustment of consumption to an unexpected change in permanent income. This assumption implies that next period’s *change* (or growth) in consumption should be unforecastable given information today.

The permanent income hypothesis also implies that there is a linear relationship between aggregate consumption, C_t ; aggregate labor income, Y_t ; and aggregate nonhuman (financial) wealth, W_t :⁶

$$(2) \quad C_t = \alpha + \beta W_t + \delta Y_t + u_t,$$

where the error term, u_t , is a discounted value of expected future (demeaned) income increases. Specifically, u_t takes the form:

$$(3) \quad u_t = \sum_{i=1}^{\infty} \rho^i (E_t \Delta Y_{t+i} - \mu),$$

where E_t is the expectation operator conditional on information available at time t , μ is the mean change in labor income, and ρ is a positive constant less than one.⁷

Equation 2 shows how modern-day consumer theory naturally implies a linear relationship between aggregate

consumption, aggregate net wealth, and aggregate labor income, much the same as in the traditional life-cycle literature with the error term given a specific interpretation. The parameters β and δ give the effect of a one-dollar increase in wealth and labor income on consumer expenditure, and can be interpreted as “marginal propensities to consume” out of wealth and income, respectively.⁸

Of course, theoretical justification is not a prerequisite for estimating a linear relationship among three variables. Nevertheless, it is helpful to have a reasonable framework with which to motivate and interpret empirical findings. Indeed, as discussed below, we find that the permanent income hypothesis—while not exactly correct—provides a reasonable approximation of much of the dynamic behavior of consumption, labor income, and wealth in U.S. time-series data. We now describe our approach to estimating the marginal propensity to consume out of wealth and labor income.

Our goal is to estimate the parameters β and δ . We begin by noting that the appropriate estimation technique will depend on the trend characteristics of the variables in equation 2. It is now widely recognized that each variable in that equation follows a stochastic trend, a fact we document in Appendix B. These trend characteristics of C , Y , and W can be described more precisely by noting that each variable appears to be nonstationary and to contain a unit root. (We refer to variables that contain a unit root as first-order integrated, or $I(1)$.) By contrast, the error term in equation 2, u_t , consists of a discounted sum of expected future *changes* in labor income. If the level of labor income is $I(1)$, the first difference of labor income will be stationary, or $I(0)$. Since u_t is simply the discounted value of these first differences, it follows that u_t will also be stationary. If consumption, labor income, and wealth are individually trending but the error term is stationary, the three variables in equation 2 must share a common trend (a unit root) while deviating from each other in the short run. In that case, we say that the variables are *cointegrated*, and the vector $\{1, -\beta, -\delta\}$ is the cointegrating vector. Appendix B presents evidence in support of the hypothesis that C , Y , and W —as measured by aggregate

time-series data—are in fact cointegrated, which implies that the error term, u_t , is stationary.

Why is cointegration important? Notice that the error term, u_t , in equation 2 will typically be both serially correlated and correlated with the regressors W_t and Y_t . In ordinary empirical applications, the effects of serial correlation are usually straightforward to address, but correlation between the error term and the regressors (regressor endogeneity) is, in practice, a much more intractable problem that can lead to inconsistent parameter estimates. By contrast, applications involving cointegrated variables have an

We find that the permanent income hypothesis—while not exactly correct—provides a reasonable approximation of much of the dynamic behavior of consumption, labor income, and wealth in U.S. time-series data.

important and unusual property: OLS estimates of cointegrating parameters (for example, of β and δ) are robust to the presence of regressor endogeneity.

To understand this result intuitively, notice that, if u_t is stationary but W_t and Y_t are individually trending, there may be some transitory correlation between W_t and u_t , or between Y_t and u_t , but the long-run correlation must be zero since trending variables must eventually diverge from stationary ones. Thus, we can exploit this property of cointegrated systems to obtain accurate estimates of β and δ using single equation estimation techniques (for example, OLS estimation) despite the fact that the regressors may be correlated with the error term.

A related implication of cointegration is that the empirical approach we employ will be robust to a variety of departures from the theory presented above. Consistent estimates of the parameters can be obtained *even if* there exist omitted explanatory variables (not accounted for by the

simple permanent income hypothesis) that are correlated with wealth and labor income. As long as the variables in equation 2 share a common stochastic trend, we can consistently estimate the parameters of that trend, circumventing many of the problems associated with identifying the influence of a change in wealth on consumption, such as how to adjust for the endogenous response of wealth to changes in economic activity or to unexpected shifts in the rate of return on financial assets.⁹

The empirical procedure discussed above relies on the presence of a single common stochastic trend—or cointegrating relationship—linking consumption, labor income, and net wealth. Consequently, the first step of our analysis is to verify that this proposition is supported in our data. As documented in Appendix B, the evidence suggests that there is a single cointegrating relationship among these variables, and we can therefore proceed to estimate the parameters of the cointegrating vector, that is, the marginal propensities β and δ .

Standard OLS estimation will produce consistent point estimates of the parameters β and δ (as long as the three variables in equation 2 are cointegrated). Nevertheless, it is important to recognize that statistical inference about the relationship among stochastically trending variables cannot be carried out using conventional standard errors. Some correction to the conventional OLS estimation method is necessary. Our approach is to use the dynamic OLS procedure of Stock and Watson (1993), which specifies a single equation taking the form

$$(4) \quad C_t = \alpha + \beta W_t + \delta Y_t + \sum_{i=-k}^k \beta_i \Delta W_{t+i} + \sum_{i=-k}^k \delta_i \Delta Y_{t+i} + u_t^*$$

where Δ is the first difference operator and u_t^* is related to

$$u_t \text{ such that } u_t^* = u_t - \sum_{i=-k}^k \beta_i \Delta w_{t+i} - \sum_{i=-k}^k \delta_i \Delta Y_{t+i}.$$

Equation 4 is estimated by OLS, but leads and lags of the first difference of the right-hand-side variables are included to eliminate the effects of regressor endogeneity on the distribution of the least squares estimator. (We also

make a correction for serial correlation of the residuals.) The coefficients on the level of wealth and labor income, β and δ , provide consistent point estimates of the marginal propensities to consume, and the corrected t -statistics we report can be compared with the standard t tables.

At first glance, equation 4 appears very similar to the traditional equation 1; both specifications include some combination of current and lagged levels of wealth and income as regressors and, in principle, the parameters β in equation 4 and b in equation 1 measure the same economic concept: the effect of a dollar increase in wealth on consumption. On closer inspection, however, it is clear that there are some important differences between these specifications. Unlike equation 1, equation 4 contains leads, in addition to lags, of the right-hand-side variables. The estimate of b from equation 1 is the sum of the coefficients on the current level and lags of the level of wealth, in order to capture the long-run impact of wealth when there are adjustment lags. By contrast, the estimate of β in equation 4 is just the coefficient on the current level of wealth, and leads and lags of the first difference are included simply to eliminate the effects of regressor endogeneity on the distribution of the least squares estimator. Similarly, equation 1 proxies for permanent income by using several lags of current income, whereas equation 4 splits permanent labor income into current labor income, which appears as a regressor, and the present discounted value of expected future labor income increases, which is subsumed in the residual term, u_t . The error term in equation 4 is specifically related to the consumer spending decision and is not assumed orthogonal to the regressors. By contrast, the error term in equation 1 is an empirical “add-on,” assumed to be orthogonal to the regressors.¹⁰

At an intuitive level, equation 4 is specified to estimate only the trend relationship linking consumption, labor income, and wealth. By contrast, equation 1, as estimated in Table 1, implicitly models both the long-run parameters and the adjustment process of consumer spending to disturbances from the equilibrium path. It is reasonable to suppose that a procedure—such as the estimation of equation 4—that separates these

two steps will produce firmer estimates of the trend component.

Before estimating equation 4, we deal with three additional specification issues that arise from the nature of the data on consumption, income, and wealth. First, note that theory typically does not rationalize distinct roles for stock market and non-stock-market wealth; total net worth enters the relationship in equation 4. Accordingly, we focus our analysis on what follows on total net worth, rather than breaking it out into stock market and non-stock-market wealth. As explained above, however, we note that quarterly fluctuations in net worth are largely driven by fluctuations in stock market wealth.

Second, standard theories of consumer behavior that imply a trend relationship linking C , Y , and W , as in equation 2, are applicable to the *flow* of consumption. Thus, durable goods expenditures should not be included in our measure of consumption since they represent replacements and additions to the asset stock, rather than the service flow from the existing stock. In what follows, we present estimates of the marginal propensities using personal consumption expenditures on nondurables and services (excluding shoes and clothing) as our expenditure measure, and we refer to this measure simply as consumption.¹¹ This consumption series is scaled up so that its sample mean matches the sample mean of total consumption expenditure, allowing a rough comparison of the size of the propensities to consume out of wealth and income estimated from these data with the size of the propensities computed from the total consumer spending series used in the first section of this article. Later, we discuss the application of these techniques to durables expenditure and how the dynamic relationship between these expenditures and wealth differs from that between wealth and the other components of consumer spending.

A final consideration is whether to express the variables in levels or in logs. In the specification above, the variables are defined in levels because we wish to estimate the effect of a *dollar* increase in wealth on consumption. Nevertheless, aggregate time-series data on consumption, wealth, and labor income appear to be closer to linear in logs than linear in levels, so heteroskedasticity is potentially impor-

tant if the regression is carried out in levels. Our solution is to use the dynamic OLS procedure above with variables expressed in logs and then to back out the implied level response using the most recent values of the consumption-income and consumption-wealth ratios. Throughout this article, we use lowercase letters to denote log variables.

Table 3 reports the results from estimating equation 4 in logs for $k = 3$.¹² Estimates are presented for the full sample period and for the sample divided into thirds. Estimated parameters are denoted with a “hat”; parameters with an l subscript give the point estimates for the log response; the implied level propensities are reported in the columns labeled “Level.”

As Table 3 shows, over the full sample period the marginal propensity to consume out of wealth, $\hat{\beta}$, is estimated to be about .046, while the marginal propensity to consume out of labor income, $\hat{\delta}$, is estimated to be about 0.72. These parameters are strongly significant according to the corrected t -statistics reported in parentheses.

Dividing the sample into thirds reveals some instability in these coefficients, echoing the findings in the first part of this article. In particular, the marginal propensity to consume out of wealth drops from about .07 in the first subperiod to somewhere between 0.025 and 0.03 in

Table 3
DYNAMIC ORDINARY LEAST SQUARES ESTIMATES
OF MARGINAL PROPENSITY TO CONSUME
OUT OF WEALTH AND LABOR INCOME

$$\text{Model: } c_t = \alpha + \beta_l w_t + \delta_l y_t + \sum_{i=-3}^3 \beta_{l,i} \Delta w_{t+i} + \sum_{i=-3}^3 \delta_{l,i} \Delta y_{t+i} + u_{l,t}^*$$

Sample Period	MPC out of Wealth		MPC out of Labor Income	
	Log	Level	Log	Level
1953:1-1997:1	0.291* (8.10)	0.046	0.605* (18.09)	0.718
1953:1-1967:4	0.380* (3.78)	0.072	0.500* (5.20)	0.615
1968:1-1982:4	0.155 (1.58)	0.031	0.729* (11.32)	0.861
1983:1-1997:1	0.151* (3.69)	0.024	0.764* (12.13)	0.907

Source: Authors' calculations.

Notes: Lowercase letters denote log values. “MPC” is the marginal propensity to consume. The sample period denotes the range of data after data points for leads and lags are removed. The t -statistics reported in parentheses are corrected non-parametrically for the effect of serial correlation.

*Significant at the 5 percent level or better.

the last two subperiods, while the marginal propensity to consume out of labor income rises from about .62 in the first subperiod to about .91 in the last two subperiods. Nonetheless, most of this instability appears to be concentrated in the early subsample, and we found that removing the post-Korean War period (by starting the sample in the first quarter of 1957 rather than in 1953) eliminated some of this instability, at least for some dynamic OLS specifications. For the post-1957 sample, the implied estimates of the marginal propensity to consume out of wealth, β , and the marginal propensity to consume out of labor income, δ , were found to be 0.04 and 0.72, respectively.

In summary, the dynamic OLS procedure employed above suggests that the estimates of the cointegrating parameters vary somewhat over time but are less unstable than those produced by the traditional methodology in the first part of this article. Moreover, much of this instability appears to be rooted in the early part of our sample. On the whole, the findings suggest that—in today’s economy—a one-dollar increase in wealth typically leads to a three-to-four-cent trend increase in consumer expenditure.

The dynamic OLS estimates of β and δ can be viewed as describing some trend relationship linking consumption, labor income, and wealth. These estimates do not tell us about the nature of short-run deviations from the trend relationship, or about the impact of quarterly fluctuations in the growth rate of wealth or labor income on future consumption growth. Such short-term dynamics are of interest, and an important property of cointegrated variables is that the cointegrating parameters may be estimated in a first-stage regression, as above, and then treated as known when estimating parameters associated with short-term dynamics (Stock 1987). We now examine the short-term relationship linking consumption, labor income, and wealth, taking into account their common trend.

SHORT-RUN DYNAMICS

We specify a model of short-run dynamics that imposes our estimated long-term trend relationship while at the same

time making allowances for the possibility of serially correlated but temporary divergences from this trend. This model takes the form

$$(5) \quad \Delta x_t = \mu + \alpha(c_{t-1} - \hat{\beta}_l w_{t-1} - \hat{\delta}_l y_{t-1}) + \sum_{j=1}^k \Gamma_j \Delta x_{t-j} + e_t,$$

where Δx_t is the vector of log first differences, $(\Delta c_t, \Delta w_t, \Delta y_t)'$, and the parameters $\hat{\beta}_l, \hat{\delta}_l$, are the previously estimated cointegrating coefficients for c_t, w_t , and y_t . The parameters μ, α , and Γ govern the short-term dynamics—that is, the relationship of consumption, wealth, and labor income growth as well as the lags of these variables and the trend deviation in the second term.¹³ Note that the parameters in this second term are the estimated coefficients from our dynamic OLS procedure.

Equation 5 is a vector autoregression (VAR) in log first differences, with the added restriction that the (log) levels of the variables share a common trend, so that last period’s deviation from trend, given by $(c_{t-1} - \hat{\beta}_l w_{t-1} - \hat{\delta}_l y_{t-1})$, is allowed to influence the current period growth of at least some of the variables. This specification is referred to as an error-correction representation, and the variable $(c_{t-1} - \hat{\beta}_l w_{t-1} - \hat{\delta}_l y_{t-1})$ as the error-correction term, since the equation takes into account any “correction” arising from last period’s deviation, or error, in the trend relationship. For any set of cointegrated variables, there is an error-correction representation, and this representation is the appropriate VAR for describing short-term dynamics among the variables in that set. An unrestricted VAR in first differences is appropriate when the variables involved are individually trending but do not have a common trend.

Table 4 summarizes the dynamic behavior of the restricted VAR in equation 5. All variables are expressed as log differences; estimates of the parameters in the error-correction term are obtained from the full post-Korean War sample using the dynamic OLS procedure discussed earlier. Results are reported for a two-lag version of the model, in accordance with findings from Akaike and Schwartz tests for lag length.

The results that appear in Table 4 are organized into three columns. For each dependent variable Δc_t , Δw_t , and Δy_t , the coefficient on the error-correction term is presented; p -values from the F -test statistic for the joint marginal significance of the block of lags of each variable and for the error-correction term are also presented in the table.

Three points about Table 4 are worth noting. First, the F -test statistics show that lagged consumption growth predicts labor income growth at the 5 percent level and growth in household net worth at the 10 percent level, but neither of the income or wealth variables predicts consumption growth. This finding is consistent with forward-looking behavior, suggesting that some consumers have information about their future asset and labor income that is not captured by lags of these variables, and that they respond to this information by changing consumption today. It also implies that an important part of the noncontemporaneous correlation between consumption and predictable changes in household net worth and labor income simply reflects the fact that consumption tends to anticipate an increase in these variables, rather than the other way around.¹⁴

Second, the F -tests in Table 4 reveal that lags of consumption growth enter significantly in the equation for consumption growth. The correlation of consumption

growth with its own lags may be the result of some adjustment delay in consumption and represents a statistical rejection of the permanent income model, which implies that the growth in consumption should be unforecastable. Nevertheless, it is clear that fluctuations in wealth do not help predict future changes in the growth of consumption once we control for lagged consumption growth.

A remaining feature of the data is that the error-correction term is significantly correlated with next period's household net worth. This finding is not predicted by the simple model discussed above. Lettau and Ludvigson (1999) develop an alternative model of forward-looking consumption behavior that allows for time-varying returns, which can account for such a correlation.¹⁵

We now move on to study the dynamic response of consumption growth to a wealth shock in order to investigate the length of time over which a change in wealth typically influences consumption growth. As a preliminary step, it is necessary to make an assumption about the timing of events, and we show the response of consumption growth to a one-standard-deviation wealth shock under two such assumptions. In the first, we assume that consumption growth may not respond to wealth within the quarter but may respond with a lag. In the second, we assume that consumption may respond to a wealth shock within the quarter. Chart 4 shows these responses for consumption growth, Δc_t . Each panel also shows two-standard-deviation error bands of these responses (dashed lines).

As the top panel shows, when we force consumption to respond with a one-period lag, a one-standard-deviation shock to the growth of wealth has virtually no impact on consumption growth at any horizon; the standard error bands are sufficiently wide that the response cannot be considered more than noise.

By contrast, the bottom panel shows the response of consumption growth to a wealth shock when we allow the former to respond contemporaneously. In this case, growth in consumption shoots up on impact, but the duration of the response is extremely short, so that by the end of the impact quarter, the effects of the shock

Table 4
ESTIMATES FROM A RESTRICTED VECTOR AUTOREGRESSION

	Equation		
	Δc_t	Δw_t	Δy_t
Joint significance of			
$\Delta c_{t-i}, i = 1 \dots 2$	0.04	0.08	0.04
$\Delta y_{t-i}, i = 1 \dots 2$	0.18	0.08	0.95
$\Delta w_{t-i}, i = 1 \dots 2$	0.61	0.08	0.39
Coefficient on error-correction term			
$c_{t-1} - \hat{\beta}_t w_{t-1} - \hat{\delta}_t y_{t-1}$	-0.001	0.476	0.113
(p -value)	(0.99)	(0.00)	(0.13)
Adjusted R ²	0.11	0.14	0.05

Source: Authors' calculations.

Notes: The table reports p -values from the F -test statistic for the joint marginal significance of the block of lags in the row for the equation with the dependent variable reported in the column. The sample period is first-quarter 1953 to fourth-quarter 1997.

are statistically negligible. This explains why wealth has virtually no impact on consumption when we force it to respond with a lag of one quarter.¹⁶

The results in the bottom panel of Chart 4 allow us to estimate the impact of stock market moves on near-term trends in consumption. The panel shows the effect of a one-standard-deviation move in wealth growth on consumption growth. A one-standard-deviation move in wealth is about 1.5 percent. The point estimate of the contemporaneous or “impact” effect on consumer spending growth


(actually, the effect in the quarter directly following the increase in wealth) of this move is about .07 percent, implying an elasticity of consumption growth to wealth growth of about .05. A \$3.5 trillion short-term movement in the stock market (comparable to those we have recently seen) equals about 10 percent of household wealth. If such a move occurred and the level of wealth then held steady (not the case recently), we estimate that there would be a .5 percent impact on consumer spending growth (2 percent at an annual rate) the next quarter. This point estimate of the impact effect is certainly interesting, but is not overwhelming in importance. Nevertheless, the great imprecision of the estimate (the two-standard-deviation error bands stretch from a negligible effect to a 1 percent effect) implies that the impact effect is quite uncertain.

What do these responses suggest for the effect of wealth changes on the *level* of consumption? Chart 4 shows that the comovement between consumption growth and an unpredictable change in wealth growth is largely contemporaneous; there do not appear to be important lagged effects in this relationship. Accordingly, when a positive wealth shock hits the economy, by the end of the impact quarter there is no further impetus from this shock to the *growth* of consumption. These responses imply that the level of consumption rises quickly to a new, permanently higher pace.

While many of the results discussed above are roughly consistent with the predictions of the permanent income hypothesis, it is clear that the permanent income interpretation is not quite right, since we know from Table 4 that consumption growth is correlated with its own lags. Serial correlation in consumption growth may be caused by any number of theoretical departures from the permanent income hypothesis, all of which can be described loosely by the umbrella term “adjustment lags.” Whatever the underlying reason for these adjustment lags, however, it appears that controlling for lags of consumption growth by itself is sufficient to account for the lags. As the bottom panel of Chart 4 illustrates, once we control for lags of consumption growth, there are no meaningful lags in the adjustment of consumption to a wealth shock.

Chart 4

RESPONSE OF CONSUMPTION GROWTH TO A WEALTH SHOCK, RESTRICTED VECTOR AUTOREGRESSION



Source: Authors' calculations.

Notes: The estimation period is first-quarter 1953 to fourth-quarter 1997. The response in the top panel is produced from a vector autoregression (VAR) for the log difference in consumption growth, labor income growth, and net worth growth, in that order. The response in the bottom panel is produced from a VAR for the same variables when consumption growth is ordered last. Both VARs impose the error-correction term. The solid lines show the response to a one-standard-deviation shock in the growth of net wealth; the dashed lines indicate two-standard-deviation error bands.

Are changes in wealth helpful in forecasting consumption growth? Put another way, should a permanent change in wealth cause us to alter our prediction of consumption growth one or more quarters ahead? We can answer this question explicitly by testing whether the specification in equation 5 improves one-quarter-ahead forecasts of consumption growth. We use a simple univariate process as a benchmark model and compare the forecasting performance of the univariate model with that of the equation 5 specification (see box).

Exploring a variety of univariate processes reveals that the log difference in consumption can be well described by a first-order autoregressive process—

although a process in which the growth of consumption is unforecastable (the log of consumption is a random walk) is not a bad approximation. The best fitting univariate processes for Δw_t and Δy_t , respectively, are a first-order moving-average process and a first-order autoregressive process. We use these univariate models below.

We make a series of one-quarter-ahead forecasts and begin by estimating each model on an initial sample period. We then make a one-quarter-ahead, out-of-sample forecast and use recursive estimation to reestimate the model, adding one quarter at a time and calculating a series of one-quarter-ahead forecasts. Forecasts are evaluated by


WHAT ABOUT DURABLES?

The results above tell us about the dynamic relationship between nondurables and services consumption and wealth. Can we characterize the short-term relationship between wealth and durables expenditure? If evidence supported the hypothesis that durables expenditure, wealth, and labor income are cointegrated, the same techniques used previously could be employed to estimate the short-term dynamics using a restricted vector autoregression (VAR) specification such as equation 5. However, the assumption of cointegration is not warranted (either empirically or theoretically) for these variables. Thus, we investigate the short-run dynamics of an unrestricted VAR in log first differences for durables expenditure, wealth, and labor income.

The chart shows the response of real durables expenditure growth to a one-standard-deviation increase in the rate of growth of net worth. Compared with the response of the scaled nondurables measure reported in the text, this response is larger in magnitude and somewhat more persistent. One quarter after the shock, durables growth increases by about 60 basis points, and the impetus to durables spending growth from this shock remains statistically positive for more than one quarter. Nevertheless, the effect on durables spending growth is only slightly more persistent than that for the nondurables measure used in the text, becoming statistically negligible by the beginning of the second quarter after a shock. By contrast, the pattern of responses for *total*

consumption (not shown) is very similar to those for the scaled nondurables consumption measure, reflecting the fact that durables expenditures represent only about 12 percent of personal consumption expenditures.

RESPONSE OF DURABLES GROWTH TO A WEALTH SHOCK, UNRESTRICTED VECTOR AUTOREGRESSION



Source: Authors' calculations.

Notes: The estimation period is first-quarter 1953 to fourth-quarter 1997. The response is produced from an unrestricted vector autoregression for the log difference in durables expenditure growth, labor income growth, and net worth growth, in that order. The solid line shows the response to a one-standard-deviation shock in the growth of net wealth; the dashed lines indicate two-standard-deviation error bands.

computing the root-mean-squared error from the set of one-quarter-ahead forecasts.

Table 5 reviews the out-of-sample forecasting performance of the restricted VAR model relative to the univariate process for each forecasted variable. Several evaluation periods are considered. First, we use a relatively long, but recent, horizon—the first quarter of 1990 to the fourth quarter of 1997. We then analyze forecast performance over four shorter, nonoverlapping horizons spanning the first quarter of 1984 to the fourth quarter of 1997. For each forecasted variable and each evaluation period, the table reports the ratio of the root-mean-squared error obtained from the univariate model to that of the restricted VAR model (equation 5). A number less than one indicates that the one-quarter-ahead forecast accuracy of the univariate process is superior to that of the VAR model.

The main features of the results may be summarized as follows. There is no evidence that the restricted VAR model consistently improves forecasts of consumption growth relative to a simple univariate process. Indeed, in four of the five evaluation periods we consider, the restricted VAR model is outperformed by a first-order autoregressive process for consumption growth. For the remaining evaluation period (the first quarter of 1987 to the fourth quarter of 1989), the two specifications perform equally well. The superiority of the autoregressive process in forecasting consumption growth is not large in magnitude. Nevertheless, the

finding that the restricted VAR model often delivers less accurate forecasts than a simple univariate model underscores the fact that quarterly fluctuations in wealth have virtually no marginal impact on future consumption growth.

These features of the results are particularly pronounced for the most recent evaluation period—the first quarter of 1994 to the fourth quarter of 1997. During this period, using the univariate model instead of a VAR model would have consistently improved forecasts of consumption growth. And, although we do not report these results in Table 5, it is worth noting that even a process in which the growth of consumption is unforecastable (the log of consumption is a random walk) would have improved one-quarter-ahead forecasts of consumption growth during this period, relative to the VAR specification.

By contrast, the VAR model appears to improve forecasts of labor income growth relative to a first-order autoregressive process for that variable: the forecasting error of the restricted VAR model for labor income growth is lower in three of the five evaluation periods we consider.

In summary, the one-quarter-ahead forecast evaluations presented in Table 5, the responses plotted in Chart 4, and the dynamic estimates displayed in Table 4 all tell the same story: Controlling for lags of consumption growth, the dynamic adjustment of consumption to an unpredictable change in wealth is largely contemporaneous, as shown by the response of consumption growth to a wealth shock in

Table 5
ONE-QUARTER-AHEAD FORECASTING PERFORMANCE OF THE VECTOR AUTOREGRESSION MODEL RELATIVE TO A UNIVARIATE MODEL

Forecasted Variable	1990:1-1997:4	1984:1-1986:4	1987:1-1989:4	1990:1-1993:4	1994:1-1997:4
Consumption growth (Δc_t) ^a	0.987	0.929	1.001	0.990	0.967
Income growth (Δy_t) ^b	1.019	1.056	0.970	0.997	1.097
Wealth growth (Δw_t) ^c	0.850	1.116	1.136	0.865	0.834

Source: Authors' calculations.

Notes: Δc_t denotes the log difference in real per capita nondurables and services consumption, excluding shoes and clothing; Δy_t denotes the log difference in real per capita after-tax labor income; Δw_t denotes the log difference in real per capita wealth. The table reports forecast evaluation statistics for predicting the variable named in the row. Each figure is the ratio of the root-mean-squared forecasting error for a univariate model relative to that of the vector autoregression (VAR); an entry of less than one indicates that the univariate model in the numerator has superior forecasting ability. Out-of-sample evaluation periods are identified in the column headings; the initial estimation period begins with the first quarter of 1953 and ends with the quarter immediately preceding the first quarter of the evaluation period.

^aThe VAR for $\Delta c_t, \Delta y_t, \Delta w_t$ is restricted (cointegration imposed); the univariate process for Δc_t is a first-order autoregressive process.

^bThe VAR for $\Delta c_t, \Delta y_t, \Delta w_t$ is restricted; the univariate process for Δy_t is a first-order autoregressive process.

^cThe VAR for $\Delta c_t, \Delta y_t, \Delta w_t$ is restricted; the univariate process for Δw_t is a first-order moving-average process.

Chart 4. Two implications of this finding are that lagged growth rates of wealth have virtually no marginal impact on current consumption growth in the restricted VAR and that a simple univariate process forecasts consumption growth as well as or better than a VAR specification, which includes wealth and labor income.

CONCLUSION

The question of how a large movement in financial wealth would affect consumer expenditure has become more pressing as fears rise that substantial market swings will cause consumer spending—and hence aggregate demand—to fluctuate sharply. In the extreme, some commentators have suggested that a prolonged downturn in stock prices could so depress consumer spending as to result in a recession (for example, see *Economist* [1998]).

How important is the stock market effect on consumption? Our results suggest that this question may be

difficult to answer partly because the trend relationship linking consumption, wealth, and labor income exhibits some instability. An important objective for future research is to investigate formally the sources and precise timing of this instability in the long-run wealth effect. Nevertheless, using a reasonable estimate of the prevailing trend relationship between wealth and consumption, we also find that the answer to this question depends on whether one is asking about today or tomorrow. Movements in the stock market today appear to influence today's consumption growth, not tomorrow's. Thus, changes in wealth in this quarter do not portend significant changes in consumption one or more quarters later. When uncertainty about the trend and impact relationship is added to the difficulties associated with wealth-based forecasts of the next quarter's consumption growth, it appears that we have a way to go before we can make inferences about movements in consumption based on movements in the stock market.

APPENDIX A: DATA

We provide a description of the data used in our empirical analysis.

CONSUMPTION

Consumption is measured as either total personal consumption expenditure or expenditure on nondurables and services, excluding shoes and clothing. The quarterly data are seasonally adjusted at annual rates, in billions of chain-weighted 1992 dollars. Our source is the U.S. Department of Commerce, Bureau of Economic Analysis.

AFTER-TAX LABOR INCOME

Labor income is defined as wages and salaries + transfer payments + other labor income - personal contributions for social insurance - taxes. Taxes are defined as $[\text{wages and salaries} / (\text{wages and salaries} + \text{proprietors' income with IVA and } C_{\text{adj}} + \text{rental income} + \text{personal dividends} + \text{personal interest income})] \times \text{personal tax and nontax payments}$, where IVA is inventory valuation and C_{adj} is capital consumption adjustments. The quarterly data are in current

dollars. Our source is the Bureau of Economic Analysis.

POPULATION

A measure of population is created by dividing real total disposable income by real per capita disposable income. Our source is the Bureau of Economic Analysis.

WEALTH

Total wealth is household net wealth in billions of current dollars. Stock market wealth includes direct household holdings, mutual fund holdings, holdings of private and public pension plans, personal trusts, and insurance companies. Our source is the Board of Governors of the Federal Reserve System.

PRICE DEFLATOR

The nominal after-tax labor income and wealth data are deflated by the personal consumption expenditure chain-type price deflator (1992=100), seasonally adjusted. Our source is the Bureau of Economic Analysis.

APPENDIX B: TESTS FOR STOCHASTIC TRENDS

This appendix describes our procedures for testing for cointegration and the results of those tests. The results for log variables are presented; results for levels of variables are very similar and are available on request. We use two types of tests: residual-based tests (designed to distinguish a system without cointegration from a system with at least one cointegrating relationship), and tests for cointegrating rank (designed to estimate the number of cointegrating relationships).

The former requires that each individual variable pass a unit-root test and is conditional on this pretesting procedure. Table B1 presents Dickey-Fuller tests for the presence of a unit root in c , y , and w over several autoregressive structures. The procedure tests the null hypothesis of a unit root against the alternative hypothesis that the series is stationary around a trend. The test statistics fall within the 95 percent confidence region and are therefore consistent with the hypothesis of a unit root in those series.

Table B2 reports statistics corresponding to the Phillips and Ouliaris (1990) residual-based cointegration tests. These tests are designed to distinguish a system without cointegration from a system with at least one cointegrating relationship. The approach applies the augmented Dickey-Fuller unit-root test to the residuals of

equation 2. Table B2 shows both the Dickey-Fuller t -statistic and the relevant 5 and 10 percent critical values.¹⁷ The hypothesis of no cointegration is rejected at the 5 percent level by the augmented Dickey-Fuller test with one or two lags, but is not rejected by the test with three or four lags. We applied the data-dependent procedure suggested in Campbell and Perron (1991) for choosing the appropriate lag length in an augmented Dickey-Fuller test. This procedure suggested that the appropriate lag length was one, implying that test results favoring cointegration should be accepted.¹⁸

Next, we consider testing procedures suggested by Johansen (1988, 1991) that allow the researcher to estimate the number of cointegrating relationships. This procedure

Table B2
PHILLIPS-OULIARIS TESTS FOR COINTEGRATION USING LOGS

Dickey-Fuller t -Statistic				Critical Values	
Lag=1	Lag=2	Lag=3	Lag=4	5 Percent Level	10 Percent Level
-4.29	-4.20	-3.75	-3.59	-3.80	-3.52

Source: Authors' calculations.

Notes: The Dickey-Fuller test statistic has been applied to the fitted residuals from the cointegrating regression of consumption on labor income and wealth. Critical values assume trending series. We use the log of consumption for non-durables and services, excluding shoes and clothing, as the dependent variable.

Table B1
DICKEY-FULLER TESTS FOR UNIT ROOTS

	Dickey-Fuller t -Statistic				Critical Values	
	Lag=1	Lag=2	Lag=3	Lag=4	5 Percent Level	10 Percent Level
Log (total wealth ^a)	-2.460	-3.067	-2.894	-3.100	3.44	3.14
Log (labor income ^a)	-0.624	-0.794	-0.829	-0.810	3.44	3.14
Log (consumption, excluding shoes and clothing ^a)	-0.363	-0.812	-0.944	-1.280	3.44	3.14

Source: Authors' calculations.

^aValues are in real per capita terms. The model includes a time trend.

presumes a p -dimensional vector autoregressive model with k lags, where p corresponds to the number of stochastic variables among which the investigator wishes to test for cointegration. For our application, $p = 3$. The Johansen procedure provides two ways of checking for cointegration. First, under the null hypothesis, H_0 , that there are exactly r cointegrating relationships, the “Trace” statistic supplies a likelihood ratio test of H_0 against the alternative, H_A , that there are p cointegrating relationships, where p is the total number of variables in the model. Second, an “L-max” statistic is formed to test the null hypothesis of r cointegrating relationships against the alternative of $r+1$ cointegrating relationships. Both of these tests for cointegration depend on the number of lags assumed in the vector error-correction structure. Table B3 presents the results obtained under a number of lag assumptions. The same effective sample (first-quarter 1954 to fourth-quarter 1997) was used to estimate the model under each lag assumption.

The critical values obtained using the Johansen approach also depend on the trend characteristics of the data. We present results for tests that allow for linear trends in the data, but we assume that the cointegrating relationship has only a constant. See Johansen (1988, 1991) for a more detailed discussion of these trend assumptions.¹⁹ The table also reports the 90 percent critical values for these statistics.²⁰

The Johansen “L-max” test results establish strong evidence of a single cointegrating relationship among the variables in equation 2. We can reject the null hypothesis of no cointegration in favor of a single cointegrating vector, and we cannot reject the null hypothesis of one cointegrating relationship against the alternative of two or three relationships across a range of trend and lag specifications. This result is also robust to every lag specification we con-

sider. While the evidence in favor of cointegration is somewhat weaker according to the “Trace” statistic (for some of the lag specifications, we cannot reject the null of no cointegration against the alternative of three cointegrating relationships), we also cannot reject the null of one cointegrating relationship against the alternative of three.

Table B3
JOHANSEN COINTEGRATION TEST: I(1) ANALYSIS WITH A LINEAR TREND IN THE DATA AND A CONSTANT IN THE COINTEGRATING RELATIONSHIP

Lag in VAR model=1					
L-max		Trace		$H_0 = r$	
Test Statistic	90 Percent CV	Test Statistic	90 Percent CV		
19.25	13.39	26.38	26.70	0	
6.14	10.60	7.13	13.31	1	
0.99	2.71	0.99	2.71	2	
Lag in VAR model=2					
L-max		Trace		$H_0 = r$	
Test Statistic	90 Percent CV	Test Statistic	90 Percent CV		
21.99	13.39	27.46	26.70	0	
4.52	10.60	5.47	13.31	1	
0.96	2.71	0.96	2.71	2	
Lag in VAR model=3					
L-max		Trace		$H_0 = r$	
Test Statistic	90 Percent CV	Test Statistic	90 Percent CV		
16.68	13.39	22.03	26.70	0	
4.55	10.60	5.35	13.31	1	
0.81	2.71	0.81	2.71	2	
Lag in VAR model=4					
L-max		Trace		$H_0 = r$	
Test Statistic	90 Percent CV	Test Statistic	90 Percent CV		
16.35	13.39	22.14	26.70	0	
4.85	10.60	5.79	13.31	1	
0.94	2.71	0.94	2.71	2	

Source: Authors' calculations.

Notes: The sample period is first-quarter 1954 to fourth-quarter 1997. Endogenous variables are the log of total wealth, the log of labor income, and the log of nondurables and services, excluding shoes and clothing. The columns labeled “Test Statistic” give the test statistic for the corresponding test above; the columns labeled “90 Percent CV” give the 90 percent confidence level for the statistics.

ENDNOTES

The authors thank Arturo Estrella, Jordi Galí, James Kahn, Kenneth Kuttner, Martin Lettau, Hamid Mebran, Patricia Mosser, Gabriel Perez Quiros, Simon Potter, Anthony Rodrigues, Daniel Thornton, Bharat Trehan, Egon Zakrajšek, and participants in the Federal Reserve System Committee Meeting in San Antonio for many valuable discussions and comments; they also thank Jeffrey Brown, Rita Chu, and Beethika Khan for able research assistance.

1. Tracy, Schneider, and Chan (1999) discuss changes in household balance sheets and the distribution of household stock holdings.
2. A good rule of thumb is that a one-point movement in the Dow Jones Industrial Average changes household sector wealth by \$1 billion to \$2 billion. The Dow gained 1,400 points from the beginning to the end of 1997, while the aggregate gain in household equity holdings during the year was about \$2.5 trillion. The Dow's low in the period since mid-1997 was 7,161 and its high through early 1999 was above 10,000—a swing of more than 3,000 points. This swing would correspond to a change in household wealth of substantially more than \$3.0 trillion.
3. Other, more recent applications of this type of estimation equation appear in Mosser (1992), Laurence H. Meyer & Associates (1994), and Poterba and Samwick (1995). Starr-McCluer (1998) examines the wealth effect using survey data.
4. Formal statistical tests of the year-to-year stability of this parameter have not been conducted. Nevertheless, the charted standard error bands give some indication of the size of the year-to-year changes relative to the statistical uncertainty of the parameter estimate.
5. Of course, this problem has been well understood for a very long time. For instance, Mishkin (1976, 1977) addressed it in a traditional life-cycle model.
6. See Galí (1990). Galí extends the infinite horizon version of the permanent income hypothesis to allow for finite horizons. Other works attempting to combine the traditional life-cycle/permanent income views with modern time-series econometrics are Blinder and Deaton (1985) and Campbell and Mankiw (1989).
7. Galí (1990) shows that ρ is a function of a constant discount rate, a constant probability of dying, and a constant rate of geometric decay in labor income growth over the lifetime.
8. Several papers have empirically tested and analyzed the permanent income hypothesis by using a single right-hand-side variable such as personal disposable income or gross national product (for example, Campbell [1987] and Cochrane [1994]). However, neither of these single measures is appropriate for our investigation because we want to estimate

the marginal propensities to consume from assets and labor income separately.

9. Our empirical approach is also robust to the possibility that consumers may have more information than the econometrician. All of the expressions that contain expectations conditional on a specific information set can be left undefined in the error term, implying that we need not make any assumption about what information consumers have in implementing our empirical procedure.
10. There are other, more subtle differences: as we explain below, consumption is defined differently in the two specifications, and equation 1 assumes a first-order autocorrelation structure for the error term requiring nonlinear estimation, while equation 4 makes a nonparametric correction for generalized serial correlation. There is also a distinction between the overall income measure used in equation 1 and the labor income measure used in equation 4. See Modigliani and Tarantelli (1975), Modigliani and Steindel (1977), and Steindel (1977, 1981) for discussions of the conceptual issues involved in using a total income measure in a traditional life-cycle/permanent income consumption model including wealth.
11. Much of the traditional literature on the life cycle also drew the same distinction between durables spending and other consumer outlays. We used total spending in Table 1 because much of the recent discussion of the wealth effect focuses on the decline in the personal saving rate. Personal saving is the difference between disposable income and all consumer outlays, including durables.
12. The results are not sensitive to choosing different values for k .
13. Klitgaard (1999) uses a similar methodology to estimate the long-run and short-run relationships of Japanese export prices to the yen exchange rate. When interest rates are not fixed but instead are time-varying, an extension of the simple model presented above implies that the *ex ante* real interest rate should be included as an additional regressor in the consumption growth equation. As is now well known, however, expected real interest rates have little impact on consumption growth. The inclusion of estimates of real interest rates in our analysis did not alter our conclusions.
14. Campbell and Mankiw (1989) document a similar result for labor income: lagged growth in nondurables and services spending is a strong forecaster of disposable personal income growth.
15. The error-correction term does not enter significantly into the equation for income growth. This latter result is inconsistent with the theory presented above since the error-correction term (equal to u_t in equation 2) comprises expected future income increases. A close

ENDNOTES (*Continued*)

Note 15 continued

examination of the labor income data reveals why: the measure of labor income we used includes transfers, the growth of which exhibits little persistence. The inclusion of transfers in labor income significantly decreases the persistence of labor income growth. Since this measure of labor income is largely unforecastable, it is uncorrelated with the lagged error-correction term. Results (not reported) show that when we use an alternate measure of labor income that excludes transfers, the error-correction term is a strong predictor of labor income growth, consistent with the theory. Nonetheless, our conclusions for consumption were not affected by our choice of labor income variable.

16. It is interesting to note that the response of wealth to its own innovation (not shown) suggests that the log of wealth is close to a random walk. This implies that the wealth shock to which consumption is responding may be viewed as a permanent increase in the level of wealth, or at least as having important permanent components.

17. Phillips and Ouliaris (1990) tabulate critical values for the augmented Dickey-Fuller t -test applied to residuals of a cointegrating equation with up to five variables.

18. An earlier version of this paper (Ludvigson and Steindel 1998) reported much less evidence in favor of cointegration because of an error in constructing the income data.

19. In choosing the appropriate trend model for our data, we were guided by both theoretical considerations and statistical criteria. Theoretical considerations implied that the long-run equilibrium relationship linking consumption, labor income, and wealth does not have deterministic trends (see equation 2), although each individual data series may have deterministic trends. Moreover, statistical criteria suggested that modeling a trend in the cointegrating relationship was not appropriate: the normalized cointegrating equation under this assumption did not correspond to any reasonable hypothesis about the long-run relationship among these variables. For example, with trends specified in the cointegrating relationship, the parameters of the cointegrating vector were often negative, an outcome at odds with any sensible model of consumer behavior.

20. The critical values are based on calculations made by Johansen and Nielsen (1993).

REFERENCES

- Ando, Albert, and Franco Modigliani.* 1963. "The Life Cycle Hypothesis of Saving: Aggregate Implications and Tests." *AMERICAN ECONOMIC REVIEW* 103, no. 1: 55-84.
- Blinder, Alan S., and Angus Deaton.* 1985. "The Time Series Consumption Function Revisited." *BROOKINGS PAPERS ON ECONOMIC ACTIVITY*, no. 2: 465-511.
- Campbell, John Y.* 1987. "Does Saving Anticipate Declining Future Labor Income? An Alternate Test of the Permanent Income Hypothesis." *ECONOMETRICA* 55, no. 2: 1249-73.
- Campbell, John Y., and N. Gregory Mankiw.* 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." NBER Working Paper no. 2924, April.
- Campbell, John Y., and Pierre Perron.* 1991. "Pitfalls and Opportunities: What Macroeconomists Should Know about Unit Roots." *NBER MACROECONOMICS ANNUAL*: 141-201.
- Cochrane, John H.* 1994. "Permanent and Transitory Components of GNP and Stock Prices." *QUARTERLY JOURNAL OF ECONOMICS* 109, no. 1: 241-65.
- Economist.* 1998. "America's Bubble Economy." April 18.
- Engle, R. F., and C. W. J. Granger.* 1987. "Cointegration and Error-Correction Representation, Estimation, and Testing." *ECONOMETRICA* 55: 251-76.
- Flavin, Marjorie.* 1981. "The Adjustment of Consumption to Changing Expectations about Future Income." *JOURNAL OF POLITICAL ECONOMY* 89, no. 5: 974-1009.
- Gali, Jordi.* 1990. "Finite Horizons, Life-Cycle Savings, and Time-Series Evidence on Consumption." *JOURNAL OF MONETARY ECONOMICS* 26, no. 3: 433-52.
- Hall, Robert E.* 1978. "Stochastic Implications of the Life-Cycle Permanent Income Hypothesis: Theory and Evidence." *JOURNAL OF POLITICAL ECONOMY* 86, no. 6: 971-87.
- Hamilton, James D.* 1994. *TIME SERIES ANALYSIS*. Princeton: Princeton University Press.
- Johansen, Soren.* 1988. "Statistical Analysis of Cointegrating Vectors." *JOURNAL OF ECONOMIC DYNAMICS AND CONTROL* 12, no. 2/3: 231-54.
- . 1991. "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models." *ECONOMETRICA* 56, no. 6: 1551-80.
- Johansen, Soren, and B. Nielsen.* 1993. "Asymptotics for Cointegrating Rank Tests in the Presence of Intervention Dummies—Manual for the Simulation Program DISCO." Manuscript, University of Copenhagen, Institute for Mathematical Statistics.
- Klitgaard, Thomas.* 1999. "Exchange Rates and Profit Margins: The Case of Japanese Exporters." Federal Reserve Bank of New York *ECONOMIC POLICY REVIEW* 5, no 1: 41-54.
- Laurence H. Meyer & Associates, Ltd.* 1994. *THE WUMM MODEL BOOK*. St. Louis.
- Lettau, Martin, and Sydney Ludvigson.* 1999. "Consumption, Aggregate Wealth, and Expected Stock Returns." Unpublished paper, Federal Reserve Bank of New York.
- Ludvigson, Sydney, and Charles Steindel.* 1998. "How Important Is the Stock Market Effect on Consumption?" Federal Reserve Bank of New York Research Paper no. 9821.
- Mankiw, N. Gregory, and Stephen P. Zeldes.* 1991. "The Consumption of Stockholders and Nonstockholders." *JOURNAL OF FINANCIAL ECONOMICS* 29, no. 1: 97-112.
- Mishkin, Frederic S.* 1976. "Illiquidity, Consumer Durable Expenditure, and Monetary Policy." *AMERICAN ECONOMIC REVIEW* 66, no. 4: 642-54.
- . 1977. "What Depressed the Consumer? The Household Balance Sheet and the 1973-75 Recession." *BROOKINGS PAPERS ON ECONOMIC ACTIVITY*, no. 1: 123-64.
- Modigliani, Franco.* 1971. "Monetary Policy and Consumption." *CONSUMER SPENDING AND MONETARY POLICY: THE LINKAGES*. Federal Reserve Bank of Boston, Conference Series no. 5.
- Modigliani, Franco, and Charles Steindel.* 1977. "Is a Tax Rebate an Effective Tool for Stabilization Policy?" *BROOKINGS PAPERS ON ECONOMIC ACTIVITY*, no. 1: 177-203.
- Modigliani, Franco, and Ezio Tarantelli.* 1975. "The Consumption Function in the Developing Economy and the Italian Experience." *AMERICAN ECONOMIC REVIEW* 65, no. 5: 825-42.

REFERENCES (*Continued*)

- Mosser, Patricia C.* 1992. "Changes in Monetary Policy Effectiveness: Evidence from Large Macroeconometric Models." Federal Reserve Bank of New York *QUARTERLY REVIEW* 17, no. 1: 36-51.
- Phillips, P., and S. Ouliaris.* 1990. "Asymptotic Properties of Residual-Based Tests for Cointegration." *ECONOMETRICA* 58, no. 1: 165-93.
- Poterba, James M., and Andrew A. Samwick.* 1995. "Stock Ownership Patterns, Stock Market Fluctuations, and Consumption." *BROOKINGS PAPERS ON ECONOMIC ACTIVITY* no. 2: 295-357.
- Starr-McCluer, Martha.* 1998. "Stock Market Wealth and Consumer Spending." Board of Governors of the Federal Reserve System, Divisions of Research and Statistics and Monetary Affairs, Working Paper no. 1998-20, April.
- Steindel, Charles.* 1977. "Personal Consumption, Property Income, and Corporate Saving." Ph.D. diss., Massachusetts Institute of Technology.
- . 1981. "The Determinants of Private Saving." In *PUBLIC POLICY AND CAPITAL FORMATION*, 101-14. Washington, D.C.: Board of Governors of the Federal Reserve System.
- Stock, James H.* 1987. "Asymptotic Properties of Least Squares Estimators of Cointegrating Vectors." *ECONOMETRICA* 55, no. 5: 113-44.
- Stock, James H., and Mark Watson.* 1993. "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems." *ECONOMETRICA* 61, no. 4: 783-820.
- Tracy, Joseph, Henry Schneider, and Sewin Chan.* 1999. "Are Stocks Overtaking Real Estate in Household Portfolios?" Federal Reserve Bank of New York *CURRENT ISSUES IN ECONOMICS AND FINANCE* 5, no. 5.

The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. The Federal Reserve Bank of New York provides no warranty, express or implied, as to the accuracy, timeliness, completeness, merchantability, or fitness for any particular purpose of any information contained in documents produced and provided by the Federal Reserve Bank of New York in any form or manner whatsoever.