

Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

**A Primer on the Economics and Time Series Econometrics
of Wealth Effects**

Morris A. Davis and Michael G. Palumbo
2001-09

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

A PRIMER ON THE ECONOMICS AND TIME SERIES ECONOMETRICS OF WEALTH EFFECTS

Morris A. Davis and Michael G. Palumbo*

Abstract

This paper reviews the statistical approach typically applied by macroeconomists to investigate the empirical link between aggregate data on household consumption, income, and wealth. In particular, we focus on studies determining whether and how much changes in net worth, such as those generated by the stock-market boom in the U.S. over the latter 1990s, are responsible for subsequent swings in the growth rate of consumer spending. We show how simple economic theory is used to motivate an econometric strategy that consists of two stages of analysis. First, regressions are used to identify trend movements shared by consumption, income, and wealth over the long run, then deviations of these series from their common long-run trends are used to help forecast consumption growth over the short run. Our discussion highlights the various judgments that researchers must make in the course of implementing this empirical approach, and we detail how specific parameter estimates describing the magnitude of the wealth effect on consumption--and even broad conclusions about its existence--are affected by making alternative choices.

January 2001

* Davis: ReturnBuy, Inc., 21641 Beaumeade Circle, Suite 319, Ashburn, VA 20147; mdavis@returnbuy.com; Palumbo: Division of Research and Statistics, Board of Governors of the Federal Reserve System, Washington, DC 20551; mpalumbo@frb.gov. We have benefitted from discussions with many of our colleagues at the Federal Reserve Board. In particular, we appreciate detailed comments and criticisms offered by Karen Dynan, Spencer Krane, David Reifschneider, and Larry Slifman. The views expressed belong to the authors and should not be attributed to the Board of Governors of the Federal Reserve System or its staff.

Introduction

The personal saving rate, as measured by the National Income and Product Accounts (NIPA), dropped from 6.5 percent at the beginning of 1995 to 0.3 percent early in 2000.¹ To give perspective on this decline, real (inflation adjusted) disposable personal income in the first quarter of 2000 was \$6.5 trillion while real personal consumption expenditures were \$6.2 trillion:² All else equal, had households maintained their propensity to spend out of income at the 1995 level, real purchases of goods and services would have been about \$300 billion – nearly 5 percent – lower than reported in the first quarter of 2000.

Figure 1 shows that at the same time the personal saving rate fell the ratio of household net worth to after-tax personal income increased dramatically. In this paper, we investigate whether there is a direct economic relationship between the rise in the wealth-income ratio and the decline in the personal saving rate. If so, then future movements in the wealth-income ratio also have implications for future household expenditures. But, if the increase in the wealth-income ratio and decline in the saving rate was mere coincidence, then the pace of household spending may be immune to changes in wealth. Understanding the link between changes in household wealth and spending – the so called “wealth effect” on consumption – is therefore critical for interpreting the recent past and considering the future.

We first will discuss how the *wealth-accumulation identity* provides an accounting link between wealth and consumption. The identity illustrates that wealth can increase through two distinct channels: People can use some of their disposable income to invest in assets (tangible or financial) instead of consuming; or, assets already owned by households (acquired through prior investments) can appreciate in price. Evaluating this identity with the macroeconomic data, one sees that the exceptional increase in wealth experienced by US households since the mid-1990s has been, in very large part, due to the rapid appreciation of equity prices over this period. Moreover, the remarkable performance of the stock market during much of the 1990s relative to

¹ In the national accounts, the personal saving rate is defined as the percentage of after-tax (disposable) income not spent by households on consumer goods and services or as interest payments to businesses. Throughout this paper we use the term “personal saving rate” and “saving rate” interchangeably.

² Both values are expressed at annual rates.

earlier periods of history suggests that in all likelihood the magnitude of the rise in equity prices was not fully anticipated by most households before the fact. As discussed below, the unanticipated nature of the increase in wealth turns out to have important macroeconomic implications.

We next will discuss how economic theory links consumption and wealth. Our basic analysis will rely on a simple benchmark model of consumer behavior known as the *life cycle model*. According to the life cycle model, households accumulate and deplete their wealth to keep their planned consumption spending roughly steady, even, for example, when their income is expected to fall as it might during retirement. In the absence of wealth “surprises,” the life cycle model predicts that wealth could vary substantially over time but that consumption spending will be relatively stable. However, if households experience an unexpected increase in their wealth, then households will formulate a new spending plan that involves a higher level of outlays indefinitely into the future. Therefore, the life cycle model suggests that predictable changes in household income and wealth (such as those reflecting new investments deliberately generated by thrift) should not lead to changes in planned spending, while household spending should respond to unexpected changes in wealth, such as from a stock-market surprise.

Economists have adapted the life cycle framework to build empirical models that quantify the relationship between aggregate consumption, income, and wealth. In the third section of the paper, we use these models to estimate the consumption response to changes in wealth. Specifically, we estimate two sets of models of consumption behavior. The first uncovers the long run relationship between consumption, income, and wealth. Under certain assumptions about the behavior of the economy, these models predict the level of consumption expected to persist after a few years in response to a change in wealth. Typical estimates of these models suggest that consumption permanently increases approximately 4 cents for each dollar that wealth increases. The dynamics of how consumption adjusts to changes in wealth is described by a second set of models that explain quarter-to-quarter, short-run movements of consumption. Estimates of these models show that, for example, in quarters when wealth increases and consumption does not immediately jump to its new long run level, spending tends to grow at an accelerated pace for several quarters in the future – until the level of consumption is brought back into line with the new level of net worth.

In the final section of the paper, we examine the sensitivity of the long-run and short-run consumption equations by examining a variety of alternative specifications. The long-run equations we investigate always imply an economically and statistically significant role for wealth in accounting for movements in consumption and saving. The estimated size of long-run wealth effects vary, with the sets of models we examine showing consumption responses between 3.0 cents and 6.5 cents for each dollar that wealth increases – suggesting that consumption increased somewhere between \$250 and \$500 billion in response to the increase in wealth from 1995 through the end of 1999. On the other hand, subtle differences in model specification seem to have large effects on estimated short-run consumption dynamics. We conclude by highlighting a few specific judgements made by different researchers that have led them to different conclusions on the nature of short-run wealth effects.

I. Wealth, Income, and Consumption in the US since 1995

As a matter of accounting, household income, consumption, and wealth must be linked together by the wealth- accumulation identity:

$$(1) \quad W_t = W_{t-1} + (Y_t - C_t) + \Delta p_t W_{t-1}.$$

The identity states that the amount of wealth owned by a household at some date (W_t) equals the level of wealth owned at the end of the previous year (W_{t-1}) plus the income saved during the year ($Y_t - C_t$), plus any capital gains or losses that might have accrued due to changes in the price of assets in the household's portfolio over the year ($\Delta p_t W_{t-1}$).³ This accounting identity holds for each household and for the aggregate across all households in the economy, as well.

³ The subscript t is used to reference the observation of a variable at a given point in time – such as the value of household wealth at the end of a particular year or quarter or the average of income and spending during the year or quarter; the subscript $t-1$ then refers to the observation in the preceding time period. The variable W refers to household net worth; that is, the market value of all assets owned by individuals minus the market value of all liabilities. The term Δp_t denotes the change in market prices for assets between periods $t-1$ and t . The variable Y represents the sum of labor income, transfer income, interest income less interest expenditures on debt, distributed capital gains from mutual funds, and dividend income (all of which are net of tax payments). Expenditures on consumer goods and services are denoted by the variable C .

In general, no individual can control the capital gains or losses on his or her existing wealth because asset prices are driven by broader market forces. However, a family can (largely) choose its level of spending and, thus, via thrift, can partially control its future wealth holdings. Accordingly, the wealth-accumulation identity allows the total increase in household wealth that occurred in any given time interval to be decomposed into contributions due to largely purposeful actions undertaken by households (saving), and those due to external market forces (net capital gains).

The level of and changes to aggregate household wealth (net worth) are estimated by the Federal Reserve Board and reported in tables B.100 and R.100 of the Flow of Funds Accounts, labeled “Balance Sheet of Households and Nonprofit Organizations” and “Change in Net Worth of Households and Nonprofit Organizations,” respectively.⁴ Selected data from these two sources are shown in table 1. We use these two Flow of Funds tables to understand the source of the rapid increase in household wealth in the second half of the 1990s. As reported in lines 1 and 2 of column 6, over the five-year period ending in 1999 the current dollar value of aggregate household net worth rose from \$24.8 trillion to \$41.8 trillion – a \$17 trillion increase (line 3).⁵ Consistent with the wealth accumulation identity of equation (1), table 1 decomposes the change in household wealth into a component for saving ($Y_t - C_t$) (labeled “Net saving”, line 4) and a separate component that measures the appreciation of existing wealth ($\Delta p_t W_{t-1}$; labeled “Holding gains”, line 5).⁶ As indicated in column 6, net saving contributed just \$1.6 trillion to the overall \$17 trillion increase in household net worth estimated over the 1995 to 1999 period. Meanwhile, the net appreciation of existing assets (holding gains) contributed \$15-1/2 trillion, or 91 percent,

⁴ At the end of the paper, we have included an appendix intended to provide the more technically inclined reader with details regarding our procedures for constructing and analyzing the data. These items are arranged by section of the paper and are alphabetized by title within each section. Refer to the item *Household sector and nonprofit organizations* in the technical appendix for additional detail.

⁵ After accounting for inflation (using the implicit deflator for personal consumption expenditures measured in chain-weighted 1996 dollars), the increase is a little less pronounced: *Real* household net worth increased from \$25-1/2 trillion to \$39 trillion over this time period.

⁶ Refer to the item *Net saving and holding gains data in the Flow of Funds* in the technical appendix for more detail. Also, the item *Net saving and the purchase of consumer durable goods* describes a technical difference between how personal saving is measured in the national income accounts and in the Flow of Funds.

Table 1: Selected Flow of Funds Data (released June 9, 2000)

			1995	1996	1997	1998	1999	1995-99
			(1)	(2)	(3)	(4)	(5)	(6)
(1)	W_t	Net worth (end of year)	27579	30098	33812	37162	41821	41821 ^a
(2)	W_{t-1}	Net worth (end of preceding year)	24814	27579	30098	33812	37162	24814 ^b
(3)		Change in net worth, (1) - (2)	2765	2519	3714	3350	4659	17007
(4)	$(Y_t - C_t)$	Net saving	350	344	244	367	323	1628
(5)	$\Delta p_t W_t$	Holding gains	2465	2114	3515	3012	4342	15448
(6)		Ratio of holding gains to total change in net worth	0.90	0.84	0.95	0.90	0.93	0.91
(7)		Memo: NIPA personal saving	302	272	271	230	156	1231

Notes: All figures are in billions of current dollars. Numbers may not sum due to rounding.

^a Net worth at the end of 1999.

^b Net worth at the beginning of 1995.

to the overall increase in net worth (lines 5 and 6).

We cannot directly observe how much of these capital gains households were expecting as they headed into 1995, but given the extraordinary performance of equity prices over the succeeding five years, it seems likely that households were largely taken by surprise. To underscore this point, we separate household wealth into “stock market wealth” and “non-stock market wealth” and then estimate annual inflation-adjusted rates of capital appreciation for these two components since 1950.⁷ Figures 2 and 3 show the historical inflation-adjusted (real) rate of capital appreciation on non-stock-market wealth and stock market wealth. In each year from 1995 through 1999, real rates of capital appreciation on both types of wealth were high, but were not abnormally large in any *particular* year. For example, the peak annual real return on stock market wealth over the 1995 - 1999 period was 31.5 percent in 1995. In comparison, households

⁷ Additional details are provided in the technical appendix items *Stock market wealth: our definition and the Flow of Funds definition* and *Calculating rates of capital appreciation for stock market and other types of wealth*.

experienced an even larger real rate of capital appreciation on equities in 1991 (34.7 percent).

However, what is remarkable about the real rates of capital appreciation experienced from 1995 through 1999 is their *persistence*, especially the ex post returns estimated for stock market wealth. Figures 4 and 5 graph the cumulative 5-year real appreciation rates to holding the two types of wealth.⁸ Figure 4 shows that from 1995 through 1999, households experienced a cumulative 5-year real appreciation of 9.6 percent on non-stock market wealth – not a typical occurrence, but one that has happened twice since the 1950s. In contrast, the sustained rate of capital appreciation on total stock market wealth experienced in the last half of the 1990s was truly exceptional (figure 5). Annually, real rates of appreciation of household stock market wealth from 1995 through 1999 were 31.5, 17.6, 25.5, 15.5, and 17.7 percent respectively, generating a 164 percent total return (including the effects of compounding) over the full 5-year period. Such a sustained pace of stock market appreciation is unmatched in postwar times. The next largest 5-year period of cumulative appreciation occurred between 1954 and 1958, when the value of households’ existing stock portfolios rose 108 percent overall.

As we have shown, most of the increase in household wealth was due to an historically large capital appreciation on existing stock market wealth. The apparent rarity of such persistent growth in equity prices makes it very unlikely for the full magnitude of the event to have been expected by households ahead of time. For example, a rough estimate of the ex-ante (say, in 1994) probability that the succeeding five years would have resulted in consecutive annual returns on stock market wealth at least as large as 15.5 percent is well below one percent.⁹ In the next section, we use some elementary economic theory to explore the implications of such a large and unexpected “wealth surprise” on household spending decisions.

II. The Life Cycle Model of Household Spending

The basic ideas and key theoretical links between consumption and wealth can be described by the life cycle model of household spending behavior, formulated by Ando and

⁸ For example, referring to Figure 4, the cumulative 5-year real appreciation on non-stock market wealth from 1989 through 1994 was *negative* 13 percent, the value recorded for 1994 on the graph.

⁹ The technical appendix item *Ex-ante probability of stock market strength from 1995 through 1999* describes this estimate.

Table 2: Numerical Example #1 of the Life Cycle Model

	Income (Y_t) (1)	Consumption (C_t) (2)	Saving out of income ($Y_t - C_t$) (3)	Wealth (end of pd.) (W_t) (4)
Youth	\$5	\$10	\$-5	\$-5
Middle-age	\$25	\$10	\$15	\$10
Old-age	\$0	\$10	\$-10	\$0

Modigliani (1963). The logic underlying the life cycle model is illustrated using the simple numerical example illustrated in table 2.¹⁰ Suppose a person lives for exactly three time periods of equal length; for convenience, call the periods “youth,” “middle-age,” and “old-age.” Further, suppose that the person works and earns income while young and middle-aged, then retires when old. Specifically, suppose she earns \$5 while young, \$25 while middle-aged, but receives no income (\$0) in old-age (column 1). Finally, assume that the young person is initially endowed with no assets and carries no liabilities, that inflation is not an issue, and that interest is not paid on savings nor accrues on debt.

In this simple world, the total amount of resources available to the person for spending throughout her lifetime equals her lifetime earnings, $\$30 = \$5 + \$25 + \0 . The person’s problem is to choose how much should be spent (and borrowed or saved) while young and middle-aged because these decisions, in turn, will determine how much she can spend during old-age. A person who simply spends all her income while young and middle-aged will have nothing to live on while retired; at the opposite extreme, a person who saves all earned income for retirement spending will be relatively miserable for most of her life.

Saying something more precise about how this consumer will allocate her spending over time, however, requires knowing much about her preferences. In postulating life cycle behavior,

¹⁰ For more formal, mathematical treatments of the life cycle model, see Deaton (1992), Muellbauer (1994), or Attanasio (1998).

Ando and Modigliani argued strongly in favor of focusing on cases in which individuals prefer to keep their levels of spending relatively steady over time. A formal treatment of this issue lies beyond the scope of this paper, but the basic principle is that most people feel that if they let consumption track their income too closely, the extra benefits from consuming “too much” in times of plenty (middle-age in our example) are outweighed by the extra costs from consuming “too little” in times of relative need (when young and old). As a result, there will be a tendency for people to want to smooth their spending relative to their income over time.

In table 2 we consider the extreme case in which the person decides to maintain *equal* spending of \$10 in each of the three periods of life. To keep the level of consumption steady in this example the person must be able to borrow when young, as well as accumulate wealth in middle-age: Here, the person borrows \$5 when young (negative saving, leaving -\$5 of wealth) to afford \$10 of consumption; she then must save \$15 in middle age in order to have accumulated, on net, \$10 in net worth to finance \$10 of spending while old.

So far, the key result to take away from this example is that according to the life cycle model people adjust their saving and wealth over time to keep their planned spending levels steady in the face of uneven income streams. Thus, according to the life cycle theory, consumption does not vary at all with *predictable* changes in wealth because people are consciously adjusting their wealth levels for the explicit purpose of keeping their consumption path smooth.

At this point, it is useful to write the level of planned spending chosen by the life cycle consumer in each period t as a multiple of the total amount of resources available to finance spending:

$$(2) \quad C_t = m_t * (H_t + W_{t-1}).$$

In the above equation, W_{t-1} denotes conventionally-defined *financial wealth* owned at the end of period $t-1$ (which is the beginning of period t); this is the same wealth variable we have considered so far, and it measures the value of financial assets and tangible property net of financial liabilities. H_t , however, is a new variable called *human wealth*, measured as current labor income Y_t plus the expected value of income to be earned in the future. Equation (2)

Table 3: Example #1 continued

	Human wealth H_t	Wealth, end of previous period W_{t-1}	Total available resources $H_t + W_{t-1}$	Consumption C_t	Propensity to consume m_t
	(1)	(2)	(3)	(4)	(5)
Youth	\$30	\$0	\$30	\$10	1/3
Middle-aged	\$25	-\$5	\$20	\$10	1/2
Old	\$0	\$10	\$10	\$10	1

represents the consumption *decision rule*; this rule states that consumption is proportional to current and expected future resources, where the factor of proportionality in a given time period, m_t , known as the *propensity to consume*, is set to keep consumption steady over the person's lifetime. At a young age, when a person's total remaining lifetime resources ($H_t + W_{t-1}$) are large, m_t will be small, and she will consume a relatively small fraction of lifetime resources. However, as a person grows older and total expected remaining lifetime resources decline, a steady consumption path requires that the propensity to spend increase.

In table 3, we reconsider the first example showing how to use equation (2) to characterize consumption choices according to the life cycle model. In the example, human wealth (H_t) is \$30 in youth, falls to \$25 in middle-age, then drops to \$0 in old-age; adding in financial wealth (W_{t-1}),¹¹ total resources available for spending (column 3) equal \$30 in youth, \$20 heading into middle-age, then \$10 in old-age. Thus, to keep consumption smooth throughout the person's lifetime, as available resources are falling with age, the propensity to consume out of total resources, m_t , must increase over time. As it turns out, in the benchmark case considered here, the propensity to consume is age-dependent – specifically, it equals one divided by the number of time periods remaining in life: 1/3 in youth; 1/2 in middle-age; 1 in

¹¹ The wealth variable in table 3, W_{t-1} , is taken from column 4 in table 2 (W_t) with a one-period lag.

old-age.¹²

A second numerical example serves to clarify the life cycle model's implications for the response of household consumption to an unexpected increase in wealth, one that is not caused by a person's own planned borrowing or saving decisions. Referring to table 4, suppose the person earns \$5 and spends \$10 while young, as in the original life cycle scenario, but then receives a surprise wealth-gift of \$20 upon entering middle-age. The person's best response, at this point, is to try to keep consumption smooth for the remaining two periods of her life. Since the person's remaining lifetime resources now equal \$40 (\$25 of income earned in middle-age minus the \$5 liability acquired while young plus the \$20 unexpected increase in wealth) she can now afford to spend \$20 while middle-aged and another \$20 while old. Thus, the extra \$10 of consumption she enjoys in both middle- and old-age (relative to the original scenario) is the "wealth effect" implied by the life cycle model of behavior. Importantly, as can be seen by multiplying $H_t + W_{t-1}$ (column 3) by m_t (column 5), this example also illustrates that the consumption decision rule – equation (2) with the age-dependent value of m_t – still applies when there is an unexpected increase in financial (or human) wealth. Finally, the example illustrates how the extra \$20 of consumption afforded by the surprise increase in wealth is used gradually,

¹² In fact, in a three-period life cycle model (without interest payments on saving or debt and without inflation) applying equation (2) with propensities to consume of 1/3, 1/2, and 1, respectively, will produce the desired steady consumption path for any income stream and any initial endowment of wealth. See the technical appendix item *Deriving the propensity to consume*.

Table 4: Example #2 (Surprise Wealth Gift)

	Human wealth H_t (1)	Wealth, end of previous period W_{t-1} (2)	Total available resources $H_t + W_{t-1}$ (3)	Consumption C_t (4)	Propensity to consume m_t (5)
Youth	\$30	\$0	\$30	\$10	1/3
Middle-aged	\$25	-\$5 + \$20	\$40	\$20	1/2
Old	\$0	\$20	\$20	\$20	1

with \$10 of extra spending occurring in each of the last two periods of the person's life.

Economists have elaborated on the simple version of the life cycle model to provide a more realistic and complex description of the process by which households make consumption and wealth-accumulation decisions. Some important extensions include: a) the possibility that it takes time for households to change their spending patterns in response to a new level of resources; b) the fact that some households may not be able to borrow as much as they would like in advance of rising income; c) the likelihood that households might want to maintain some additional assets as a precaution against possible bad luck given the uncertainty surrounding future, unpredictable economic outcomes; d) the ability of households to earn interest on their accumulated assets, and the requirement of making interest payments on debts. These extensions of the simple model deliver the same basic results – namely, that as long as households can borrow against future expected earnings or anticipated increases in future wealth, they will attempt to keep a relatively smooth consumption path throughout their lifetime.¹³ However, these extensions also imply some room for empirical deviations from the strict behavioral predictions made by the simple theory. In particular, they allow for the possibility that current spending may be sensitive to changes in income, even when those movements are

¹³ In the technical appendix, the item *Realistic versions of the life cycle model* mentions some extensions of the benchmark model that allow it to mimic reality better, but leave the basic results intact. The technical appendix item *Borrowing constraints* mentions a potentially important case in which the basic results fail to hold: Namely, if some households cannot readily borrow at prevailing interest rates, then consumption will respond to changes in income and wealth that were predictable in advance.

predictable in advance, and that household spending may respond slowly to permanent changes in income and wealth. They also suggest that aggregate consumption may be related to past movements in interest rates, unemployment, and other variables that help to predict future movements in income or proxy for uncertainty.

Summarizing, expected and unexpected changes in wealth or income have different implications for household spending according to the simple life cycle model. Any anticipated change in wealth or income will have no direct causal effect on consumption because it would have been appropriately incorporated into the desired level of spending when first recognized to occur, rather than after the expected event. By definition, unexpected events cannot be planned for in advance, so households optimally respond to them after the fact: For example, because surprise increases in wealth or income generate higher levels of lifetime resources, once they occur they will generate higher levels of spending over the remainder of the household's planning period.

Life cycle theory enables us to qualitatively assess the influence of the gains in wealth over the latter half of the 1990s on spending behavior. Had the increase in household net worth over the period largely come from increased personal saving, theory predicts there would not have been a burst in consumption – life cycle consumers save in order to keep their spending smooth. However, as we have shown, the increase in net worth was probably driven by unexpected increases in equity prices, which life cycle theory predicts would result in a sustained jump in household spending.

But, by how much? According to the life cycle model, knowing the magnitude of wealth effects on consumption comes down to knowing the value of the propensity to consume, m_c , for the economy as a whole. Many factors, such as the profile of life expectancies, real interest rates, borrowing constraints, and households' aversion to risk, influence the value of m_c . Within the context of the strict life cycle theory, some analysts have tried to guess the value of m_c given data on interest rates and life expectancies.¹⁴ Most researchers, however, attempt to estimate m_c

¹⁴ Poterba (2000) reports a range of wealth-effects estimates derived by evaluating the theoretical decision rule similar to equation (2) that is derived from a somewhat more complex life cycle model. The range reflects alternative assumptions for interest rates and life expectancies. Poterba's calculations for the propensity to consume out of a dollar's increase in wealth range from 2.7 to 10.3 cents (refer to table 3 in his paper).

using statistical relationships between observed spending and household resources. In the next section, we describe the econometric methods and macroeconomic time series data typically used to produce such estimates of the propensity to consume.

III. Estimating Wealth Effects

So far, we have provided evidence that households most likely did not expect the run-up in stock prices from 1995-1999 before the fact, and have demonstrated that unexpected increases in wealth should have led to some increase in household spending according to the life cycle model of consumption. We now address the task of using macroeconomic data to quantify how much consumption might have increased in response to the sharp rise in wealth during the second half of the 1990s.¹⁵

A. *Wealth Effects over the Long Run*

For convenience, we rewrite the consumption decision rule derived in the previous section for an arbitrary household (labeled household “i”):

$$(3) \quad C_{t,i}^* = m_{t,i} * (H_{t,i} + W_{t-1,i}).$$

In equation (3) $C_{t,i}^*$ is planned consumption as predicted by the life cycle model, $m_{t,i}$ represents the age-dependent propensity to consume out of total resources; as before, total resources are defined as the sum of human wealth ($H_{t,i}$) and financial wealth ($W_{t-1,i}$).

Adding up equation (3) across all households in the economy yields an aggregate “consumption function” of the form:

$$(4) \quad C_t^* = b * (H_t + W_{t-1}),$$

where C_t^* is aggregate planned consumption, H_t is aggregate human wealth (current and

¹⁵ A complementary approach estimates wealth effects from household-level, or “microeconomic,” data. In a recent paper, Dynan and Maki (2000) summarize the extant research using this approach and present new evidence.

expected future income) and W_{t-1} is aggregate household net worth.¹⁶ In equation (4) we have replaced the household and age-specific propensity to consume, $m_{t,i}$, with a fixed aggregate propensity to consume, b . This transformation is innocuous if the proportion of the economy's total wealth held by the different age-cohorts of the population is relatively stable over time.

Two substantial impediments exist for using equation (4) directly in empirical research. First, expectations of future income are unobservable; consequently, the value of H_t is not available in any data. To account for this, researchers often use a proxy for human wealth and recognize that they will only be estimating an approximate version of equation (4). Second, as mentioned above, households may not always behave exactly as predicted by the simplest version of the life cycle theory. For example, changes in C_t could trail those in C_t^* because consumers may adjust their spending with a lag in response to news about their income or wealth, as would be the case if it takes time to decide what items should be purchased (like a car or vacation) in response to good news. Thus, actual consumption, C_t , and planned consumption C_t^* may not always equal each other, even if human wealth could be perfectly measured.¹⁷

Denoting the proxy measure of human wealth as $\hat{H}_t$, the empirical counterpart to equation (4) looks like

$$(5) \quad \begin{aligned} C_t &= \hat{C}_t + \epsilon_t \\ &= b * (\hat{H}_t + W_{t-1}) + \epsilon_t, \end{aligned}$$

where C_t is actual consumption and $\hat{C}_t$ is target consumption as predicted by the approximation to the life cycle model using the proxy for human wealth instead of actual (but unobservable) H_t .¹⁸ In the parlance of econometricians, the difference between actual consumption (C_t) and the proxy for target consumption ($\hat{C}_t$) is often called an *error* or *gap*, and is denoted in equation (5)

¹⁶ Deaton (1992, pp. 45-7) derives the aggregate life cycle consumption function explicitly.

¹⁷ In addition, as described in the next section, departures in household behavior from the strict predictions of the simplest life cycle theory mean that the short-run response of spending when income or wealth change might not correspond exactly to equation (4).

¹⁸ Note the difference between C_t^* and $\hat{C}_t$. C_t^* is based on actual (but unobserved) human wealth, H_t , while $\hat{C}_t$ is based on the human wealth proxy, $\hat{H}_t$.

by ϵ_t . (For convenience, we will drop the use of the term “proxy” and simply refer to $\hat{C}_t$ as target or planned consumption.) As we discuss in the next section of the paper, if actual consumption is temporarily above or below target consumption, then knowledge of the current period’s value for the gap (ϵ_t) may be useful in predicting next period’s consumer spending.

As an aside, note that equation (5) implies that, on average, consumption should be a constant proportion of total resources, but that, in general, consumption will not be a constant proportion of financial wealth (or income) alone. Rather, according to (5), C_t / W_{t-1} should fall if W_{t-1} increases – unless H_t increases proportionally.¹⁹ Thus, the ratio of consumption to financial wealth will mimic any trend in the ratio of human to financial wealth, which could be expected to fall in the case of a sudden jump in equity prices. Nonetheless, some commentators have interpreted the recent decline in the ratio of consumption to financial wealth as indicating a relatively “mild” wealth effect operating over the latter half of the 1990’s, even though, by itself, observations on C / W are uninformative about the behavioral response of consumption to changes in wealth.

Researchers have differed in their approaches to finding a good proxy for human wealth. Perhaps the most common approach is to assume that human wealth is proportional to current income, or $\hat{H}_t = kY_t$;²⁰ in this case equation (5) becomes

$$(6) \quad \begin{aligned} C_t &= \hat{C}_t + \epsilon_t \\ &= aY_t + bW_{t-1} + \epsilon_t, \end{aligned}$$

where $\hat{C}_t = aY_t + bW_{t-1}$ and a equals b times the factor of proportionality, k .

For statistical reasons relating to the trending nature of aggregate income, wealth, and spending data, econometric analysis cannot be directly applied to equation (6); in particular, ordinary regression techniques are flawed because the average magnitude of the error term ϵ_t is

¹⁹ To see this, divide both sides of equation (5) by W_{t-1} : $C_t / W_{t-1} = b + b(H_t / W_{t-1}) + \epsilon_t / W_{t-1}$.

²⁰ Macklem (1994) and the FRB/US model (Brayton and Tinsley, 1996) specify and estimate an econometric model for human wealth rather than assuming that expected income is a constant proportion of current income. See the technical appendix item *Proportionality of human wealth to current income: household and aggregate data* for further discussion of this issue.

not stable over time – a property referred to as nonstationarity.²¹ There are two common ways to circumvent this statistical problem. First, the variables in equation (6) can be re-scaled by dividing the terms in equation (4) by the level of disposable income, as in:

$$(7) \quad \begin{aligned} \frac{C_t}{Y_t} &= \frac{\hat{C}_t^1}{Y_t} + \frac{\epsilon_t^1}{Y_t} \\ &= a + b \frac{W_{t-1}}{Y_t} + \frac{\epsilon_t^1}{Y_t}. \end{aligned}$$

Second, some economists apply a logarithmic approximation to equation (6), which takes the form:

$$(8) \quad \begin{aligned} \log(C_t) &= \log(\hat{C}_t^2) + \epsilon_t^2 \\ &= c_0 + \alpha \log(Y_t) + \beta \log(W_{t-1}) + \epsilon_t^2, \end{aligned}$$

where c_0 is a constant and α and β relate the natural logarithms of income and wealth to the natural logarithm of target life cycle spending, $\log(\hat{C}_t^2)$.²² The transformed series for consumption, income, and wealth that appear in equations (7) and (8) are *cointegrated* – that is, the transformed errors, ϵ_t^1/Y_t and ϵ_t^2 , are stationary. Consequently, the coefficients can be estimated using ordinary least squares regression techniques.

We will refer to equation (6) as *life cycle model 1* – which we estimate using the cointegrated form given by equation (7) – and to equation (8) as *life cycle model 2*. To begin,

²¹ Specifically, by relating the levels of consumption, income, and wealth – three macroeconomic series that exhibit very strong upward trends over time – the absolute value of the error in equation (7) tends to grow over time. This property of nonstationarity precludes direct estimation of the coefficients using ordinary least squares, which assumes that the average magnitude of the error is approximately constant over time. The solutions to this problem that we describe in the text transform the data in a way that does not remove the strong trends from the macroeconomic variables, but which produces an error term that has no trend. Technically speaking, our procedures produce a *cointegrating* relationship among the transformed consumption, income, and wealth series. An introduction to this econometric issue is provided in the technical appendix item *Cointegration of consumption, income, and wealth data*.

²² Campbell and Deaton (1989) discuss the use of the logarithmic approximation.

our baseline estimates use “standard” publically available measures of consumption, income, and wealth; specifically, we use NIPA aggregate personal consumption expenditures as our measure of spending, NIPA personal disposable and transfer income as our measures of household income, and total household net worth from the Flow of Funds.²³ In a later section, we investigate how the estimates of long-run wealth effects are affected by making alternative data choices.

Before estimating the equations, we make an additional change: Economists sometimes subdivide income into its constituent components. One subdivision is to parse income into the component derived from government transfer programs, Y_t^{τ} , and all other income ($Y_t - Y_t^{\tau}$). Splitting income in this manner is often justified by noting that transfer income tends to be received by the elderly or poor, who should have a higher propensity to consume than the rest of the population; for example, in 1999, social security and Medicare benefits account for more than 80 percent of NIPA transfer income (Y_t^{τ}).²⁴ For life cycle model 1, this generalization yields²⁵

$$(9) \quad \frac{C_t}{Y_t} = a_1 + (a_2 - a_1) \frac{Y_t^{\tau}}{Y_t} + b \frac{W_{t-1}}{Y_t} + \frac{\epsilon_t^1}{Y_t}.$$

We estimate equations (8) and (9) using the Stock-Watson procedure over the period from 1960:Q1 to 2000:Q1, and report the coefficients in table 5.²⁶ The statistically significant

²³ The consumption, income, and wealth series are measured in real – inflation adjusted – chain-weighted dollars based on prices in 1996.

²⁴ Some researchers also prefer to subdivide wealth into its components. For example, Muellbauer (1994) argues for splitting household net worth into separate categories for relatively liquid and illiquid assets, on the grounds that differences among assets in liquidity and the distribution of ownership could imply different aggregate propensities to consume. Below we show how life cycle consumption plans and the implied wealth effects are affected by separately including stock market and non-stock market wealth in the equations.

²⁵ Because the logarithm of transfer income does not turn out to add significant predictive power to equation (8), we omit this variable from life cycle model 2.

²⁶ The technical appendix item *Stock-Watson procedure for estimating cointegration coefficients* contains a brief description of the statistical technique.

coefficient estimates of model 1 imply that planned consumption increases 68 cents when non-transfer disposable income increases by one dollar (the value of a_1 transformed from dollars to cents), 89 cents when transfer income increases by one dollar (the transformed value of a_2), and 3.9 cents whenever wealth increases by one dollar (the transformed value of b). The value of b therefore directly tells us the estimated wealth effect for model 1: 3.9 cents to the dollar.

Interpreting the coefficients for model 2 is less straightforward because it is specified in logarithmic form, which links *percentage* movements in consumption to *percentage* movements in income and wealth.²⁷ Thus, a “cents to the dollar” wealth effect estimate from model 2 depends on the levels of consumption and wealth, which, of course, vary over time. Using the average values of consumption and wealth during the last half of the 1990s generates a wealth effect from model 2 of about 3.3 cents to the dollar, only a bit lower than the estimate of 3.9 cents to the dollar of model 1.

Even though the models predict somewhat different relationships between movements in consumption, income, and wealth, both models produce nearly identical and close fits to the aggregate data. Figure 6 shows the predicted level of consumption arising from both models over the full sample period, while figure 7 highlights the predicted levels from the two models over the most recent five-year period (1995 - 2000).²⁸ Figure 8 presents the errors from models 1 and 2 in percent (to make them comparable over time), and shows that the “typical” miss is 1.1

²⁷ The technical appendix item *Interpretation of the coefficients on wealth in models 1 and 2* derives this result.

²⁸ For details on the construction of this series, refer to the technical appendix item *Predicted value of the level of consumption from the logarithms model (model 2)*.

Table 5: Estimates of Models 1 and 2, 1960:Q1 to 2000:Q1

	Variable (1)	Coefficient (2)	Estimate (3)	t-statistic (4)
Model 1	$Y_t - Y_t^\tau$	a_1	0.68	15.86
	Y_t^τ	a_2	0.89	2.36
	W_{t-1}	b	0.039	4.83
Model 2	constant	c_0	-0.58	-5.93
	$\log(Y_t)$	α	0.83	17.90
	$\log(W_{t-1})$	β	0.19	4.28

Notes: Both models are estimated using the Stock-Watson procedure in which 2 leads, 2 lags, and the contemporaneous first difference of the estimating equation's variables are included as stationary regressors. t-statistics are calculated using the procedure in p. 611 of Hamilton (1994). The t-statistic for the coefficient a_2 tests if $a_2 = a_1$; all other t-statistics test for nonzero coefficients.

percent.²⁹

Even though the t-statistics for b and β both imply statistical significance, it is reasonable to ask how much the wealth variables improve the fit of these equations. To check, we set b and β in equations (8) and (9) to zero, re-estimate the income parameters of both life cycle models, and calculate new target values of consumption. The size of the “typical” percentage consumption error (the standard deviation of the prediction errors) from the modified models is nearly double that of the original specifications.³⁰ Thus, these calculations confirm an important role for household wealth in explaining aggregate consumption in the post-war U.S.

We next use the coefficients from the estimated life cycle equations to compute the

²⁹ Figure 8 plots the percentage errors from the models, calculated as 100 times $(1 - \hat{C}_t/C_t)$. Here, we use the term “typical” in reference to the estimated standard deviation of the percentage errors to the level of consumption, which is 1.06 percent for life cycle model 1 and 1.08 percent for model 2.

³⁰ Excluding wealth from the life cycle models in this way increases the standard deviation of the estimated errors to about 2.0 percent of consumption from the 1.1 percent values reported in footnote 30.

predicted level of real consumption in 2000:Q1 assuming that household net worth had grown at a 4 percent annual rate from 1995 through 1999 rather than the actual rate of 8 percent per year; we also assume that income keeps to its historical path in this counterfactual exercise.³¹ Under this alternative, real household wealth in 2000:Q1 would be \$8 trillion less than in the historical data. Comparing simulations under the alternative and historical wealth paths, therefore, provides quantitative estimates from the two models regarding the size of “wealth effects” over the latter half of the last decade.

Table 6 contains the results of this experiment. Column (2) reports fitted values for consumption from the models in 1995:Q1; this is the value of real consumption that is consistent with the realized levels of income and wealth in 1995:Q1 according to the estimated equations. Column (3) shows for each of the models the predicted values for consumption that are consistent with the level of income and wealth actually realized in 2000:Q1; column (4) reports predicted values in 2000:Q1 under the alternative assumption for wealth. Thus, the rightmost column of table 6 presents estimates from the two models for wealth effects over the last half of the 1990s relative to our counterfactual assumption: Model 1 implies wealth effects of \$308 billion; according to model 2, the effects are a bit smaller at \$258 billion. In terms of consumption growth, both models imply that wealth effects added one percentage point to annual consumption growth – from 3.4 percent to 4.4 percent per year over average – over the

³¹ This experiment approximately holds the wealth-income ratio constant at its level in 1995:Q1.

Table 6: Wealth Effect on Consumption between 1995:Q1 and 2000:Q1

Model	Fitted C, 1995:Q1	Fitted C, 2000:Q1	Predicted C, alternate wealth path 2000:Q1	“Wealth Effect” (3) - (4)
(1)	(2)	(3)	(4)	(5)
1	4942	6132	5824	308
2	4945	6113	5855	258

1995:Q1 - 2000:Q1 period.³²

These estimates of wealth effects closely bracket the \$300 billion counterfactual level of consumption we introduced at the beginning of the paper. By implication, according to these simple life cycle regression models, the exceptional increase in household net worth relative to its long-run average pace can explain essentially all of the drop in the personal saving rate observed over the past five years, given realized income growth.

The wealth effect estimates we presented in table 6 represent long-run effects. As we describe in a later section, the exact size of the estimated long-run effects are somewhat sensitive to alternative specifications of the life cycle equations. However, before presenting an analysis of these forms of sensitivity, we first describe the approach used by researchers to model the quarter-to-quarter dynamics of consumption and, thus, measure wealth effects over the short run.

B. Wealth Effects over the Short Run

The relatively small errors shown in figure 6 suggest that empirical life cycle models like 1 and 2 pin down the effects of income and wealth on consumption over the long run with a reasonable degree of accuracy. That said, figure 8 also shows that sizable short-lived prediction errors do occur, but tend to be quickly reversed. In light of this pattern, Davidson et al (1979), asked whether consumption might move in the current period to actively “correct” a previous error. If consumption exhibits this tendency, referred to as *error correction*, then short-run

³² The estimates and simulations of these models excluding data since 1994 are very similar to those produced using the full sample. Therefore, our estimates of wealth effects over the latter half of the 1990s are not significantly influenced by the most recent data.

wealth effects could look quite different than those estimated to occur over longer periods of time. Tinsley (1993) formally demonstrates that an extended life cycle theory, which incorporates slowly changing household spending habits and adjustment lags in the response to economic news, implies that an equation like (5) approximately describes the relationship between consumption, income, and wealth over long periods of time, but also implies that spending exhibits error correction-type behavior in the short run.

More pertinent to our empirical analysis, however, is the following point: The regression coefficients presented in table 5 can be interpreted as representing the effects of changes in income and wealth on consumption only if it can be demonstrated that spending levels move in response to movements in income and wealth, rather than the coefficients exclusively reflecting adjustments in income or wealth that occur in response to changes in consumption. The process of error correction in consumption growth implies that such a channel running from changes in wealth to changes in consumption does exist. This section presents an investigation of error-correction behavior in consumption growth over the short run.

A stylized example of error correction. Using a stylized numerical example, table 7 demonstrates how error-correction behavior can produce different short-run wealth effects compared with the case in which household spending immediately and fully adjusts to changes in income and wealth. In the example, spending in the economy starts out at \$9600 (row 4 column 2), which equals the level of planned consumption (row 4 column 2) as predicted by the life cycle equation (7).³³

Suppose that a rise in equity prices in year 1 suddenly lifts household net worth from \$50,000 to \$60,400 (row 1 column 3), while income remains at \$10,000. The life cycle equation suggests an immediate increase in the level of *planned* consumption to \$10,016 from the initial \$9600 level (row 3). However, in this example we assume that because of adjustment lags, households actually increase their spending in year 1 only to \$9850 in response to the wealth gains (row 4). The actual increase in spending thus lags the planned increase and, as a result, an

³³ This example is based on an equation for the level of planned consumption that is very close to the estimates for model 1 from table 5: $C_t^* = 200 + 0.7(Y_t - Y_t^e) + 0.9Y_t^e + 0.04W_t$. In this simple example, non-transfer income, $Y_t - Y_t^e$, and transfer income are set equal to 8000 and 2000, respectively, for all time periods. Also, interest is assumed not to accrue to wealth over time, but wealth does accumulate due to each year's saving.

Table 7: Wealth Effect with “Slow” Adjustment

		Notation	Start	Year 1	Year 2	Year 3
		(1)	(2)	(3)	(4)	(5)
(1)	Wealth	W_{t-1}	50,000	60,400	60,554	60,617
(2)	Income	Y_t	10,000	10,000	10,000	10,000
(3)	Target consumption	C_t^*	9600	10,016	10,022	10,025
(4)	Actual consumption	C_t	9600	9850	9940	9980
(5)	Percentage error	$100*(1-C_t^*/C_t)$	0%	-1%	-1/2%	-1/4%
(6)	Cumulative wealth effect			\$230	\$315	\$355

error on the order of 1 percent of the level of consumption occurs in year 1 (row 5, column 3). Continuing, no changes in income or wealth are realized in year 2 (rows 1 and 2 of column 4), so that target spending remains near the year-1 level--at \$10,022 (row 3)--but households take the opportunity to partially close the 1-percent gap between their actual and planned levels of spending: Actual spending increases to \$9940 in year 2, but still lies 1/2 percent below the target.³⁴ In year 3, actual consumption moves closer yet to planned consumption and the size of the error falls to 1/4 percent.

In this example, households deliberately but gradually move their actual consumption toward target consumption over time. As this occurs, the measured wealth effect – shown in row 6 – gradually rises. If we let the example play out further, the measured increase in actual consumption would approach \$400 – the product of the assumed propensity to consume out of wealth, 0.04, and the unanticipated change in net worth, \$10,000, which comes directly from equation (7). Therefore, the short-run pattern of wealth effects, but not the amount that emerges over the long run, is governed by the speed with which households adjust their actual spending to match the change in planned consumption warranted by movements in income and wealth.

³⁴ For simplicity of exposition, we assume that households close one-half of last period’s error in the current period. Later in this section, we present empirical estimates of the speed of error correction for household spending in the U.S.

Figure 9 compares the consumption paths resulting from the sudden jump in wealth shown in table 7 for four different speeds of adjustment. C(I) (the solid line) graphs the consumption response given in table 7. As the line shows, consumption immediately, but not fully, adjusts to the wealth shock and then slowly closes the error that opened up. In contrast, C(II) (the dashed line) plots the course of consumption if consumption immediately and fully adjusted to the wealth shock. C(III) (the dot-dash line) plots the course of consumption if consumption at first partially adjusts as in C(I), but then more quickly closes the error. Comparisons of C(I), C(II), and C(III) show that the exact size of the wealth effects in the first few years depends on the speed of households' reaction, but the long-run effects are identical. If households react immediately and fully, as shown by path C(II), then no error correction type adjustment is necessary since there is no error to correct. In this case, long-run wealth effects exactly equal short-run wealth effects. In contrast, if households do not fully react immediately (cases C(I) and C(III)), then wealth effects in the first few periods will be of lesser magnitude than long-run effects, but the error correction process eventually will bring actual spending into line with the long-run prediction of the life cycle model.

The line labeled C(IV) shows one possible path of consumption if consumption does not error correct. In this case, all of the wealth effects on consumption operate in year one, and the 2% gap that opens up in year 1 remains as long as income and wealth remain at their year-1 levels. Thus, consumption never adjusts to the change in wealth. This highlights the fact that one needs to find evidence that spending exhibits error-correction behavior in the short run to assert that changes in income or wealth eventually generate changes in consumption in the long run (all else being equal). In other words, one would not necessarily care to interpret the regression coefficients b or β in table 5 as representing a long-run "wealth effect" on household spending without also citing evidence of error correction-type behavior in the aggregate consumption data.

Of course, there is a problem with the gap never closing as in C(IV): It flies in the face of the actual behavior of consumption, income, and wealth. Figure 8 already has shown the strong tendency for life cycle errors to revert fairly quickly to their long-run average value of zero. If consumption does not error correct, then the only way that the life cycle errors can get back down to zero is for either income or wealth (in some combination) to fall in the future.

This result would lower the planned level of spending and eliminate the gap between actual and planned consumption. However, the situation in which wealth and income error-correct but consumption does not runs counter to some simple macroeconomic intuition.³⁵ For example, an implication of income error correcting is that a rise in the stock market (unaccompanied by an appreciable rise in spending) would forecast a future drop in income, which is inconsistent with the tendency for stock prices to lead economic activity. Alternatively, if wealth error corrected, then increases in wealth today caused by a runup in equity prices would forecast a future decline in equity prices. This scenario is problematic because, it leaves open a reasonably easy way to make money in the stock market.

Estimating the speed of error correction. The short run consumption dynamics described by table 7 and figure 9 were intentionally kept simple to provide a transparent illustration of error correction behavior. In contrast, in the real world households can receive news about their income and wealth prospects at any time. This implies that the consumption target is always moving, so that the macroeconomic data will never show the consumption error smoothly returning to zero, as it does in our artificial examples based on just one sudden jump in wealth and a steady level of income. Accordingly, estimating the error correction speed requires additional econometric analysis.

Researchers draw conclusions about whether consumption moves to correct recent errors by estimating whether the lagged error from a life cycle equation – for example, ϵ_{t-1}^2 from equation (9) – is a significant predictor of consumption in period t, given last period's consumption, as in:

$$(10) \quad \log(C_t) = \log(C_{t-1}) + \gamma_1 + \gamma_2 \epsilon_{t-1}^2 + \gamma_3' x_t + v_t ,$$

where γ_1 is a constant term, ϵ_{t-1}^2 is a measure of last period's gap between actual and target

³⁵ Note that evidence of error correction in household spending does not rule out error correction in income or wealth, nor does evidence of the latter tendencies rule out the former. Quantifying all of these various tendencies requires developing full econometric models for the short run behavior of household income and wealth, which would be applied along with the formulation for aggregate consumption. Here, we do not make an attempt at this more ambitious research project, which has been pursued, for example, by Brayton and Tinsley (1996), Ludvigson and Steindel (1999), and Kiley (2000).

consumption, and x denotes other variables that researchers have found to influence the short run behavior of consumption. These variables often include real interest rates, unemployment rates, measures of consumer sentiment, and lagged growth rates of consumption, income, and wealth. Their inclusion is typically motivated by the extensions to the simple life cycle model that we alluded to above.³⁶ For example, the unemployment rate or consumer sentiment indexes are intended to capture the precautionary behavior of households, while including variables that predict the growth rate of income can proxy for the potential effects of borrowing constraints on consumer spending. Meanwhile, including lagged growth rates of consumption, income, and wealth help to capture additional short-run dynamics in the reactions of these variables to transitory shocks that do not affect the target level of consumption (and, thus, do not involve error correction, per se). Note that to the extent that these variables could be affected by movements in the stock market – or could include lagged changes in wealth directly – they offer an additional channel for wealth to affect consumer spending in the short run.

Equation (10) can be rewritten as:

$$(11) \quad \Delta \log(C_t) = \gamma_1 + \gamma_2 \epsilon_{t-1}^2 + \gamma_3' x_t + v_t ,$$

where the left hand side variable now represents the growth rate of consumer spending.^{37,38}

The interpretation of equation (11) is that if the coefficient γ_2 takes a significantly negative value, then we can say that current-period consumption moves to correct an error

³⁶ See Muellbauer (1994) for a more specific explanation as to why these macroeconomic variables might influence consumption growth over the short run.

³⁷ If the growth rate of consumption is “small,” then it is approximately equal to the difference in the logarithm. See the technical appendix item *Log-difference of consumption as an approximation to the growth rate* for details.

³⁸ Analogously, we investigate whether the lagged error from life cycle equation (8), ϵ_{t-1}^1 , is a significant predictor of changes in the ratio of consumption to income, as in:

$$\Delta(C_t/Y_t) = \delta_1 + \delta_2 \epsilon_{t-1}^1 + \delta_3' x_t + v_t .$$

leftover from last period (ϵ^2_{t-1}). Under error correction, γ_2 is negative because periods in which actual consumption lies *below* target consumption and consumption errors are *negative* ($\epsilon^2_{t-1} < 0$) are followed by periods with *faster*-than-usual consumption growth to close the gap. The larger (in absolute value) is the value estimated for γ_2 , the more quickly households appear to adjust their consumption behavior to correct errors that opened up in earlier periods, and the more quickly the full life cycle wealth effect would be phased in. Hence, the reference in the research literature to γ_2 as the *error correction speed* for consumption growth. A finding that γ_2 is essentially zero would imply that current consumption growth is unaffected by the size of last period's error, meaning that whatever effect wealth is going to have on consumption tends to show up immediately in the data (as in example C(IV) in figure 9), or operates through the other variables included in the short run equation.

Note that finding a statistically significant error correction speed for consumption is sufficient for concluding that sustained changes in income and wealth tend to result in changes in spending eventually, but it does not necessarily mean that actual consumption will be noticeably affected every time the stock market moves. According to equation (11), quarterly volatility in wealth levels caused by stock market gyrations would show up as erratic consumption data only to the extent that the error correction speed, γ_2 , takes a value close to (negative) one. The range of estimates we report below are consistent with households adjusting their spending gradually in response to changes in their wealth and income levels. The gradual response speed – in combination with the relatively small propensity to consume out of wealth (3 to 6 cents-to-the-dollar) – has two important implications. First, only relatively long-lived movements in resource levels – such as those brought on by the 5-year stock market boom in the latter 1990s – can be expected to have a noticeable effect on consumption; second, only a small fraction of any one quarter's variation in wealth can be expected to actually show up in the aggregate consumption data.

Table 8 reports a group of estimated error correction speeds using consumption errors constructed from life cycle models 1 and 2 and alternative specifications for the other variables (x_t) included in the short-run consumption growth equation (10). Specification A does not include any x variables in the error correction equation. Specification B includes lagged growth rates of consumption and wealth in x , and to these specification C adds predicted current income

growth, the lagged change in the unemployment rate, and the lagged real federal funds interest rate to x .³⁹

The upper row in the table reports the one-quarter response from model 1 of the change in the spending-income ratio to a one percent life cycle consumption error. All of the estimates in this row are statistically significant. The estimates indicate that periods when the level of consumption lies below the target level (as measured using life cycle model 1 and the realized levels of income and wealth) are followed by significant upward movements in the ratio of consumption to income. Put another way, periods of consumption shortfalls are followed by quarters in which consumption grows significantly faster than disposable income.⁴⁰ As can be seen, a negative percentage point life cycle consumption error this quarter leads to a 0.21 percentage point increase in the ratio of consumption-to-income next quarter according to specification (A), and a 0.15 percentage point increase in the ratio according to specifications (B) and (C). These latter speed-of-adjustment estimates imply that roughly half of any consumption gap is closed within four quarters, while the former implies that about 60 percent is

³⁹ The prediction of current income growth is derived from a regression of income growth on lagged growth rates of consumption, income, and wealth, and lagged unemployment and real interest rates.

⁴⁰ The technical appendix item *Accounting for changes in the consumption-income ratio* relates a change in the consumption-income ratio to the difference between growth rates of consumption and income.

Table 8: Estimates of Error Correction Speeds for Consumption

Dependent Variable	Life cycle consumption error	Estimated Error Correction Speed (and diagnostic statistics)		
		Specification (A)	Specification (B)	Specification (C)
$\Delta(C_t/Y_t)$	Model 1: $\epsilon^1_{t-1} / Y_{t-1}$	-.21	-.15	-.15
		t-stat. = -3.63	t-stat. = -2.52	t-stat. = -2.14
		R ² = .08	R ² = .16	R ² = .28
$\Delta \log(C_t)$	Model 2: ϵ^2_{t-1}	-.024	-.045	-.13
		t-stat. = -.44	t-stat. = -.88	t-stat. = -2.39
		R ² = .00	R ² = .14	R ² = .45

Notes: Specification (A) includes only an intercept and the lagged life cycle consumption error as regressors. Specification (B) adds to (A) lagged growth rates of consumption and wealth (lagged income growth was not statistically significant). Specification (C) adds to (A) instrumented income growth, lagged consumption growth, the lagged change in the unemployment rate, the inflation-adjusted Federal Funds interest rate, and the current level of consumer sentiment.

complete after a year.⁴¹

The lower row in table 8 reports estimated error correction speeds in short-run equations that use the consumption growth rate $\Delta \log(C_t)$ as the dependent variable. Specifications (A) and (B) present weak evidence that household spending grows in the current period simply to correct past life cycle consumption errors, as the estimated error correction speeds are not significantly different from zero. Specification (C), however, indicates that when a negative percentage point consumption error opens up this quarter, consumption tends to grow an additional 1/2 percentage point at an annual rate next quarter to partially close the gap.⁴² Thus, according to specifications (A) and (B), future movements in income and wealth are largely responsible for closing consumption gaps that might open up in the current quarter; in contrast, specification (C) implies

⁴¹ All else being equal, the proportion of the gap closed after four quarters is roughly $1-(1+\gamma_2)^4$, which is 0.48 for $\gamma_2 = -0.15$ and 0.61 for $\gamma_2 = -0.21$.

⁴² The error correction speeds in the lower row of table 8 reflect quarterly growth rates; multiplying by 4 (approximately) converts the implied consumption growth to an annual rate (0.5 percentage point $\approx 4*0.13$).

that households indeed boost consumption growth in the short run to allow their actual spending to catch up to the new target level. Note, however, that specification (C) provides a considerably better explanation for consumption growth observed over the sample period compared with specifications (A) and (B), as indicated by the much larger R^2 statistic; the additional macroeconomic series included in specification (C) contribute most to this equation's relatively good in-sample fit, although the statistically significant error correction speed helps importantly as well.

The range of coefficient estimates reported in table 8 suggest that the quantitative importance of error correction for near-term consumption growth and, hence, wealth effects over the short run are sensitive to model and data specification. Next, we examine wealth effects with an explicit focus on the sensitivity of estimates.

IV. Sensitivity of Estimated Wealth Effects

In a recent article, Poterba (2000) highlights the existence of disparate estimates of wealth effects in the research literature. In this section we extend the analysis by showing how estimates of long-run and short-run wealth effects are affected by various differences in the specification of empirical models as well as differences in exactly which data are used. Our analysis identifies a range for the magnitude of wealth effects over the short and long run, and suggests how particular judgmental decisions can lead researchers to obtain results of differing magnitudes.

A. Sensitivity of Long-Run Effects

Sensitivity to sample period. Other authors have documented that the estimates of a given long-run model's coefficients are sensitive to the sample period chosen for estimation; see, for example, Ludvigson and Steindel (1999) and Poterba (2000). We do not focus on this issue because, in our view, researchers are best served by estimating the life cycle equations using the longest span of data available: Due to the properties of cointegration, the long-run statistical relationships between consumption, income, and wealth are better estimated when a longer span of time series data are used.

Sensitivity to model and data specification. To start, we introduce two variants of the life cycle models investigated in the previous section. In the first, labeled (1a) and (2a) in table

Table 9: Models Used to Highlight Long-run Uncertainty

Model	Equation
1	$C_t = a_1(Y_t - Y_t^r) + a_2 Y_t^r + b W_{t-1} + \epsilon_t^1$
1a	$C_t = a_1(\tilde{Y}_t - Y_t^r) + a_2 Y_t^r + b W_{t-1} + \epsilon_t^{1a}$
1b	$C_t = a_1(\tilde{Y}_t - Y_t^r) + a_2 Y_t^r + b_1 W_{t-1}^s + b_2(W_{t-1} - W_{t-1}^s) + \epsilon_t^{1b}$
2	$\log(C_t) = c_0 + \alpha \log(Y_t) + \beta \log(W_{t-1}) + \epsilon_t^2$
2a	$\log(C_t) = c_0 + \alpha \log(\tilde{Y}_t) + \beta \log(W_{t-1}) + \epsilon_t^{2a}$
2b	$\log(C_t) = c_0 + \alpha \log(\tilde{Y}_t) + \beta_1 \log(W_{t-1}^s) + \beta_2 \log(W_{t-1} - W_{t-1}^s) + \epsilon_t^{2b}$

9, we simply remove property income from disposable income.⁴³ This new definition of income, denoted $\tilde{Y}_t$, therefore only includes after-tax labor and transfer income. The motivation behind this change is that, according to the life cycle theory, property income equals the return earned on financial wealth, and so should not be included in the proxy for human wealth. Including property income in Y_t potentially muddles the distinction between the estimated propensities to consume out of income and wealth.⁴⁴

In the second variant, labeled (1b) and (2b) in table 9, we continue to use an income series that excludes property income (as in variants (1a) and (2a)), but also split total household net worth into the value of stock market wealth, W_{t-1}^s , and the value of all other assets and liabilities, $W_{t-1} - W_{t-1}^s$. Some researchers have stressed the merits of distinguishing between more and less liquid components of net worth to better fit the aggregate consumption data (for example, Muellbauer (1994)). This variant allows stock market effects on consumption to be larger or smaller than the effects of appreciation in house prices or other asset prices, as might be

⁴³ Property income is corporate dividends, net interest, rental income, and proprietors income.

⁴⁴ Note that by this line of reasoning, property income should not be held “constant” when wealth adopts a different path, as in our counterfactual calculations of table 6. However, in the data, property income – such as dividends and interest – does not move in lock step with household net worth, somewhat mitigating the force of this issue. As mentioned, we used disposable personal income as measured in the national accounts in the baseline models of the previous section in the interest of generating easily replicable results.

the case if ownership of stocks were concentrated among groups that had different preferences or faced different constraints than the groups primarily owning other assets, or if the risk and return characteristics of different assets were important determinants of the propensity to consume.

We estimate the coefficients of equations (2), (2a), and (2b) directly and, for the statistical reasons discussed above, estimate the coefficients of models (1), (1a), and (1b) using the “ratio” form analogous to equation (8). Coefficient estimates appear in table 10. Comparing the estimate of b (column (3)) in models (1) and (1a) shows that using the income series $\tilde{Y}$ boosts our estimate of the long-run response of consumption to a dollar increase in wealth from 3.9 cents to 6.3 cents; similarly for logarithmic equations (models (2) and (2a)), after making the appropriate calculations, the long-run response of consumption to a dollar increase in wealth increases from 3.3 to 5.9 cents when we exclude property income from disposable income.⁴⁵ Comparing models (1a) and (2a) with their counterparts (1b) and (2b) reveals smaller estimated coefficients for stock market wealth, column (3), than for other types of wealth, column (4). According to model (1b), an increase in stock market wealth of one dollar leads to a 5.7 cent increase in consumption over the long run whereas model (1a) shows a 6.3 cent-to-the-dollar effect from an increase in total net worth. For the equations in logarithms, we calculate that propensity to consume out of stock market wealth in model (2b) is 3.8 cents-to-the-dollar, compared with the 5.9 cent propensity to consume out of total net worth in model (2a).^{46,47}

How important is the sensitivity of long-run estimates? Column (5) translates the differences between model variants in terms of their implied wealth effects over the recent

⁴⁵ In models (2) and (2a) we calculate the cents-to-the-dollar wealth effect as β times the average value of the consumption-to-wealth ratio from 1995 through 1999, 0.173. See the technical appendix item *Interpretation of the coefficient on wealth in models 1 and 2*.

⁴⁶ Recall, because model (2b) is in logarithm form, the cents-to-the dollar wealth effect equals β_1 times the average value of the consumption-to-stock market wealth ratio over the 1995-99 period, 0.55.

⁴⁷ For reference, note that, using their full sample, Ludvigson and Steindel (1999) report a propensity to consume out of total wealth of 4.6 cents-to-the-dollar, while Kiley (2000) reports an analogous estimate of 3.3 cents. Focusing more specifically on the consumption effects of a change in stock market wealth, Brayton and Tinsley (1996) suggest a propensity to consume of 3.0 cents, while Laurence Meyer and Associates (1994) suggest 4.2 cents.

sample period.⁴⁸ The \$301 billion stock market wealth effect for the recent period implied by model (2b) is markedly smaller than the \$460 billion estimate that derives from model (2a), which restricts stock market and other wealth to influence consumption symmetrically, but is still larger than the \$258 billion effect estimated earlier from model (2) (which did not separate out property income). By contrast, subdividing net worth does not result in widely divergent long-run wealth effects according to life-cycle model variants (1a) and (1b) (\$499 billion vs. \$457 billion); and both of these estimates also are much larger than when property income is not separated from other forms of income.

We think the range of estimates shown in table 10 is representative of results that other researchers could be expected to find using other variants of these equations. All of these variants imply an important role for wealth in explaining consumption behavior over the long run. Even though the exact size of the estimated long run wealth effect varies across the life

⁴⁸ The counterfactual path for total wealth used to construct the recent wealth effects in table 10 matches that used to construct the estimates in table 6. For models (1b) and (2b) - which separate stock market and non-stock market wealth - our counterfactual simulations for the table 10 estimates attribute nearly all of the total increase in net worth over the 1995-99 period (relative to historical growth rates) to stock market wealth. This simulation is consistent with our earlier analysis of stock market and non-stock market returns over the latter 1990s.

Table 10: Estimates of Various Models, 1960:Q1 - 2000:Q1

Model	a_1	a_2	b or b_1	b_2	Wealth Effects, 1995 - 1999
	(1)	(2)	(3)	(4)	(5)
1	0.68 (15.86)	0.89 (2.36)	0.039 (4.83)		308
1a	0.73 (9.06)	1.04 (2.11)	0.063 (5.41)		499
1b	0.67 (7.72)	0.94 (2.55)	0.057 (5.52)	0.080 (5.20)	457
	c_0	α	β or β_1	β_2	
2	-0.58 (-5.93)	0.83 (17.90)	0.19 (4.28)		258
2a	-0.68 (-4.82)	0.69 (11.40)	0.34 (5.98)		460
2b	-0.69 (-6.06)	0.61 (10.42)	0.07 (7.95)	0.36 (6.53)	301

Notes: All models are estimated using the Stock-Watson procedure in which 2 leads, 2 lags, and the contemporaneous first difference of the estimating equation's variables are included as stationary regressors. T-statistics are listed in parentheses and are calculated using the procedure in p. 611 of Hamilton (1994). The t-statistic for the coefficient a_2 tests if $a_2 = a_1$; all other t-statistics test for nonzero coefficients.

cycle variants, the existence of this effect is never called into question.

B. Sensitivity of Short-Run Effects

Sensitivity to the specification of the long-run consumption equation. As noted earlier, the speed of error correction (γ_2) is estimated by regressing the growth rate of consumption on the lagged percentage consumption error (and other control variables). Because different empirical variants of the life cycle model produce different quarterly paths for consumption errors, estimates of γ_2 , and hence the short-run response of consumption to movements in income and wealth, will differ depending on which empirical representation of the life cycle model is used. In practice, we have found that estimated error corrections speeds from the three life cycle model variants presented above (using otherwise identical versions of

equation (10)) often differ by 4 or 5 percentage points, which is a substantial proportion of the typical-sized estimate of 10 to 15 percent.

To demonstrate this, it is instructive to examine the percentage consumption errors arising from life cycle model variants over the recent period. Figure 10 plots the percent consumption errors arising from models 2, 2a, and 2b from 1994:1 to 2000:1.⁴⁹ From 1994 through 1997, the models produce roughly the same time series for life cycle consumption errors, indicating that actual consumption exceeded the life cycle target level by about 2 percent in early 1994, but that the gap was essentially erased by early 1997. However, the three life cycle model variants tell appreciably different stories about consumption behavior between 1997 and 2000. Model 2 suggests that actual consumer spending rose closely in line with increases in the target level from early 1997 through the first half of 1999; after that, however, target spending predicted by the model grew relatively slowly, leaving actual spending about 1-1/2 percent above the target at the stock-market peak in 2000:Q1. In contrast, target spending predicted by variant 2a – the variant with the largest wealth effect, see table 10 – grew very rapidly from 1997 to 2000. As a result, actual spending fell below the target during 1997, and the gap remained largely unclosed through 2000 despite robust growth in actual spending. Meanwhile, life cycle variant 2b takes a middle position between these other two, and suggests that spending in 2000:Q1 was close to its target.

The logic of error correction for consumer spending means that uncertainty about last period's life cycle consumption error, which reflects uncertainty about the appropriate long-run model, directly translates into uncertainty about consumption growth over the near term. If, for example, consumption in 2000:Q1 was "too high," as suggested by models 2 and 2b, then – all else equal – error correction means that consumption growth should have slowed down considerably in 2000:Q2. In contrast, if the level of spending was "too low" in 2000:Q1, as indicated by model 2a, then error correction would have dictated rapid consumption growth in the second quarter of 2000. In fact, real PCE grew at an annual rate of 3.1 percent in 2000:Q2,

⁴⁹ We do not plot the errors from models 1, 1a, and 1b in figure 10 because they are very similar to their model-2 counterparts. Figure 11 shows the estimated consumption errors for all six variants of the life cycle equations over the full sample period. See the technical appendix item *Consumption errors from the six variants of the life cycle equations* for a discussion of figure 11.

compared with 7.6 percent in 2000:Q1 and an average rate of 4.9 percent over the 1997-1999 period.

Sensitivity to the specification of the short-run consumption equation. Estimated error correction speeds also are affected by the specification of the short-run equation (11), essentially without regard to which of the six variants is used to model the long-run behavior of consumption, income, and wealth. Below, we report our most interesting findings.⁵⁰

Wealth information used to construct the life cycle error (ϵ_{t-1}). There is a disparity in the way previous researchers have dated household net worth when constructing the life cycle errors used in the short run equations for consumption growth. This specification issue arises because wealth is a stock variable and, thus, is measured at a point in time, while consumption is a flow variable and, thus, is measured over an interval of time. In empirical practice, it is necessary to take a stand about the amount of resources perceived by households to be available to finance their spending over the relevant interval. In predicting consumption growth in 2000:Q1, for example, some researchers measure the life cycle error using net worth valued at December 31, 1999, while others use net worth valued at September 30, 1999. Comparing the two columns in table 11 shows that basing life cycle errors on more recent wealth information increases estimated error correction speeds for consumption.⁵¹ In the context of explaining consumer spending, accounting for the most recent developments in household income and wealth seems intuitively appealing. Furthermore, as is evident from comparing R^2 statistics between the two columns of table 11, using the more up-to-date information improves the fit of the short-run equations quite a bit – particularly in the equations with narrow specifications for x .

Information set used to predict income growth. It is common for researchers to include a variable in equation (11) that captures the movements in consumption growth that are directly linked to expected movements in current income growth, as would be dictated by extended life cycle model with borrowing constraints or precautionary behavior on the part of households.

⁵⁰ Additional detail regarding each of the following specification issues is presented in the technical appendix item *Factors affecting estimated error-correction speeds for consumption*.

⁵¹ Note that judgments about the timing of the measurement of net worth does not appreciably affect estimates of the coefficients in the long run equations; hence, this issue is not important for quantifying long run wealth effects.

We investigated how two different approaches to this modeling issue affected estimates of error correction speeds. Comparing the upper two rows of table 11 reveals that predicting income growth using a broad range of macroeconomic information (the middle row) yields substantially larger error correction speeds than using a forecast based only on the previous quarter's income growth (the upper row). Naturally, using a more broad information set provides a better fit to the historical income data; as it turns out, the better income prediction leads to an improvement in the fit of the short-run equation for consumption growth – indicated by the larger R^2 statistics reported in the middle row of table 11.

Other variables included in the short run consumption equation. As the bottom row of table 11 suggests, we have found that including additional macroeconomic factors – the change in the unemployment rate, the inflation-adjusted Federal Funds interest rate, and the level of consumer sentiment – improves the ability of the short run equation to explain historical consumption growth; it also assigns a sizable role to error correction in the short run consumption dynamics.

How important is the sensitivity of short-run estimates? A measure of the practical importance of uncertainty regarding short-run wealth effects is given by the product of two factors: the size of a typical long-run consumption error and the degree of uncertainty about the true error correction speed for consumption growth. Using the standard deviation of consumption errors from the life cycle equations we have estimated as a guide, the typical (absolute) value of the long run gap between actual and planned consumption is about 1 percent. Thus, following a period with an average sized consumption error, an estimate for the error correction speed (γ_2) of 12 percent would produce a forecast for quarterly consumption growth that is about 1/2 percentage point (annual rate) faster than the forecast from a model with no error correction.⁵² This compares with the 3-1/2 percent average annual growth rate for

⁵² The calculation is: $(1.0 \text{ percent} * (0.12 - 0.00)) * 4 \approx 0.5 \text{ percentage point at an annual rate.}$

Table 11: Sensitivity of Estimated Error Correction Speeds for Consumption

Information Used to Forecast Income Growth	Wealth Information Used to Construct t-1 Life Cycle Error			
	End of Period $t-2$ (1)		End of Period $t-1$ (2)	
Lagged income growth ^{a.}	-.019		-.068	
	t-stat = -.43	R ² = .11	t-stat = -1.65	R ² = .21
Broad information set ^{b.}	-.059		-.095	
	t-stat = -1.42	R ² = .23	t-stat = -2.46	R ² = .34
Specification (C) of table 8 ^{c.}	-.070		-.12	
	t-stat = -1.91	R ² = .44	t-stat = -3.36	R ² = .47

Notes: Total personal consumption expenditures is the measure of household spending and life cycle consumption errors are constructed using variant 2a. In the upper two rows, current consumption growth is regressed on the lagged life cycle error, lagged consumption growth, lagged wealth growth, and the forecast of income growth.

a. One lag of income growth is the only variable used to forecast current income growth.

b. Four lags of income, wealth, and consumption growth, the unemployment rate, and the real federal funds interest rate are used to forecast current income growth.

c. The broad information set is used to forecast current income growth. This equation also includes lags of the change in the unemployment rate, inflation-adjusted Federal Funds interest rate, and the current level of consumer sentiment as additional explanatory variables for consumption growth.

consumer expenditures over the 1960:Q1-2000:Q1 period.

Summing up. The short-run equations consistently reveal error correction in the ratio of consumption to income, indicating that a sudden increase in wealth that is not fully accommodated by a simultaneous increase in spending (as predicted by the long-run equation) results in a period in which consumption grows faster than income to close the gap. Taken together, the evidence also leans toward a period of faster than normal consumption growth (irrespective of income growth) following such a sudden increase in wealth. While we report a few estimated error correction speeds that turn out to be small and statistically insignificant, a general result is that empirical analyses that take a broader macroeconomic perspective in terms of the data employed in the model produce sizable and statistically significant error correction

speeds, and provide a better fit to the historical data on consumption growth.

V. Conclusions

The sharp runup in equity prices experienced in U.S. markets since the end of 1994 shows up clearly in the official estimates of the value of household net worth. Basic economic theory suggests that, unless consumers anticipated the full extent of the bull market ahead of time, spending should have been boosted by the gains in wealth. Indeed, we have shown that the aggregate data are qualitatively consistent with the basic theoretical prediction. The econometric models that we have estimated based on the long-run relationships between consumption, income, and wealth are consistent with a wealth effect in the range of 3 to 6 cents-to-the dollar, depending on particular model specification. These estimates translate into additional consumer spending from 1995 through 1999 of around \$250 billion to \$500 billion compared with what would have been expected under more typical appreciation of stock prices.

To interpret the long-run regression coefficients as reflecting how spending adjusts to changes in income and wealth – rather than operating the other way around – we have investigated the extent to which quarterly consumption growth acts to correct errors that open up between actual spending and its long run target level. In this respect, our empirical results are somewhat less certain, as different specifications lead to noticeable differences in the estimated behavior of consumption to a change in wealth over the short run. That said, we find that the better-fitting equations we examined tend to show error correction behavior for household spending that is consistent with direct wealth effects operating over the short run. The range of error-correction speeds imply that households adjust their spending only gradually upon realizing gains (or losses) in their income and wealth levels. Thus, only increases in income or wealth that are sustained for a sufficiently long period of time – such as those generated by the extraordinary stock market performance over the 1995-1999 period – can be expected to noticeably affect levels of consumption.

In summary, our analysis demonstrates that the appreciable increase in consumer spending in the U.S. over the latter half of the 1990s aligns quite well with the exceptional performance of the stock market based on estimated historical relationships among the aggregate data on consumption, income, and wealth.

References

- Ando, Albert, and Franco Modigliani. "The 'Life Cycle' Hypothesis of Saving: Aggregate Implications and Tests," *American Economic Review*, vol. 53 (March 1963), pp. 55-84.
- Attanasio, Orazio P. "Consumption Demand," NBER working paper no. 6466 (March 1998).
- Bliner, Alan S., and Angus Deaton. "The Time Series Consumption Function Revisited," *Brookings Papers on Economic Activity*, vol. (1985), pp. 465-521.
- Brayton, Flint, and Peter Tinsley. "A Guide to FRB/US: A Macroeconomic Model of the United States," Federal Reserve Board Financial and Economic Discussion Series Working Paper no. 1996-42 (October 1996).
- Campbell, John, and Angus Deaton. "Why is Consumption So Smooth?," *Review of Economic Studies*, vol. 56 (July 1989), pp. 357-373.
- Carroll, Christopher D., and Lawrence H. Summers. "Consumption Growth Parallels Income Growth: Some New Evidence," in B. Douglas Bernheim and John B. Shoven (eds.), *National Saving and Economic Performance*, Chicago: Chicago University Press, 1991, pp. 305-43.
- Davidson, James E. H, David F. Hendry, Frank Srba, and Stephen Yeo. "Econometric Modeling of the Time-Series Relationship between Consumers' Expenditure and Income in the United Kingdom," *Economic Journal*, vol. 88 (December 1978), pp. 661-92.
- Deaton, Angus. *Understanding Consumption*. Oxford, U.K.: Oxford University Press, 1992.
- Dynan, Karen, and Dean Maki. "Does Stock Market Wealth Matter for Consumption?" mimeo, Federal Reserve Board (July 2000).
- Flavin, Marjorie. "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy*, vol. 89 (December 1981), pp. 974-1009.
- Hamilton, James D. *Time Series Analysis*. Princeton, N.J.: Princeton University Press, 1994.
- Kiley, Michael T. "Identifying the Effect of Stock Market Wealth on Consumption: Pitfalls and New Evidence," mimeo, Federal Reserve Board (July 2000).
- Laurence H. Meyer and Associates. *The WUMM Model Book*. St. Louis: L.H. Meyer and Associates, 1994.

- Ludvigson, Sydney, and Charles Steindel. "How Important is the Stock Market Effect on Consumption?" Federal Reserve Bank of New York *Economic Policy Review*, vol. 5 (July 1999), pp. 29-51.
- Macklem, Tiff. "Wealth, Disposable Income and Consumption: Some Evidence for Canada." Bank of Canada Working Paper (November 1994).
- Modigliani, Franco, and Alberto Ando. "The Life Cycle Hypothesis of Saving: Aggregate Implications and Tests," *American Economic Review*, vol. 53 (March 1963), pp. 55-84.
- Muellbauer, John. "The Assessment: Consumer Expenditure," *Oxford Review of Economic Policy*, vol. 10 (Summer 1994), pp. 1-41.
- Muellbauer, John, and Ralph Lattimore. "The Consumption Function: A Theoretical and Empirical Overview," in H. Pesaran and M. Wickens (eds), *Handbook of Applied Econometrics*, Oxford, U.K.: Blackwell, 1995, pp. 221-311.
- Poterba, James M. "Stock Market Wealth and Consumption," *Journal of Economic Perspectives*, vol. 13 (Spring 2000), pp. 99-118.
- Tinsley, Peter A. "Fitting both Data and Theories: Polynomial Adjustment Costs and Error-correction Decision Rules," Federal Reserve Board Financial and Economic Discussion Series Working Paper no. 1993-21 (1993).

Some International Estimates of Consumption Functions

- Aron, Janine, and John Muellbauer. "Financial Liberalisation, Consumption and Debt in South Africa," mimeo, Nuffeld College, University of Oxford, U.K. (October 1999).
- Beenstock, Michael, Yaakov Lavi, and Akiva Offenbacher, "A Macroeconomic Model for Israel 1962-1990: A Market Equilibrium Approach to Aggregate Demand and Supply." Bank of Israel, Research Department, Discussion Paper No. 92.07, April 1992.
- Berg, Lennart and Reinhold Bergstrom, "Housing and Financial Wealth, Financial Deregulation and Consumption-The Swedish Case." *Scandinavian Journal of Economics*, Vol. 97 No. 3, p. 421-439, 1995.
- Black, Richard and Vincenzo, et. al., "Forecasting and Policy System: The Core Model." Reserve Bank of New Zealand, Research paper No. 43, August 1997.

- Bradley, John, Karl Whelan, and Jonathan Wright, "HERMIN Ireland." *Economic Modelling*, Vol. 12 No. 3, pp. 249-274, 1995.
- Brodin, P. Anders and Ragnar Nymoen, "Wealth Effects and Exogeneity: The Norwegian Consumption Function 1966(1) - 1989(4)." *Oxford Bulletin of Economics and Statistics*, Vol. 54 No. 3, p. 431-454, 1992.
- Brooks, Ray and Darren Gibbs, "A Model of the New Zealand Economy Reserve Bank Model XII." *Economic Modelling*, Vol. 11 No. 1, pp. 5-86, 1994.
- Campos, Julia and Neil R. Ericsson, "Constructive Data Mining: Modeling Consumers' Expenditure in Venezuela." International Finance Discussion Papers, No. 663, April 2000.
- Christensen, Anders Moller and Dan Knudsen, "MONA: A Quarterly Model of the Danish Economy." *Economic Modelling*, Vol. 9 No. 1, pp. 10-74, 1992.
- FKSEC: A Macro-Econometric Model for the Netherlands. Central Planning Bureau, 1992.
- Gardner, Robert, "Consumption and Saving in the TRYM Model," Australian Treasury Working Paper (no date given).
- Hendry, David F. "HUS Revisited," *Oxford Review of Economic Policy*, vol. 10 (Summer 1994), pp. 86-106.
- Jacobs, Jan and Elmer Sterken, "The IBS-CCSO Quarterly Model of the Netherlands: Specification, Simulation, and Analysis." *Economic Modelling*, Vol. 12 No. 2, pp. 111-163, 1995.
- Lattimore, Ralph, "Australian Consumption and Saving." *Oxford Review of Economic Policy*, Vol. 10 No. 2, 1994.
- MORKMON: A Quarterly Model of the Netherlands for Macro-Economic Policy Analysis. Monetary Monographs, No. 2. The Netherlands Bank, 1985.
- Pretorius, C.J. and S. Knox, "Private Consumption Expenditure in the Macro-Econometric Model of the Reserve Bank." Quarterly Bulletin, South African Reserve Bank, June 1995.
- Rae, David, "NBNZ-DEMONZ: A Dynamic Equilibrium Model of New Zealand." *Economic Modelling*, Vol. 13 No. 1, pp. 91-165, 1996.

Rossi, Nicola and Ignazio Visco, "Private Saving and Government Deficit in Italy (1951-1990)."
Bank of Italy Working Paper No. 178, October 1992.

Rossi, Nicola and F. Schiantarelli, "Modelling Consumers' Expenditure, Italy 1965-1977."
European Economic Review, Vol. 17, p. 371-391, 1982.

Technical Appendix

This appendix is intended to provide the more technically inclined reader with details regarding our procedures for constructing and analyzing the data. The items in the appendix are arranged by the section of the paper referencing them, and are alphabetized by title within each section.

I. Wealth, Income, and Consumption in the US since 1995

Calculating rates of capital appreciation for stock market and other types of wealth

To calculate nominal rates of capital appreciation on wealth, we assume that all net saving measured in the Flow of Funds occurred in non-stock market assets. We use this assumption because in the flow of funds accounts essentially all accumulation occurs in non-stock market assets.⁵³ To convert all nominal rates to real rates of capital appreciation, we use the rate of price inflation of NIPA consumer expenditures.

Ex-ante probability of stock market strength from 1995 - 1999

Using our definition of real stock market capital appreciation, we calculate that annual real stock returns are distributed with a mean of 7.75 percent and a standard deviation of 16.0 percent. (These returns are consistent with an alternative estimate based on CRSP data, which show a real value-weighted return on stocks of approximately 8 percent per year.) We assume that annual stock returns are identically and independently distributed from the Normal distribution. Given our estimate of the mean and variance of annual returns, the probability of experiencing a rate of return of 15.5 percent or larger in any given year is 31.5 percent. The ex-ante probability of 5 consecutive annual real returns of 15.5 percent or larger is estimated as $(0.315)^5 \approx 0.003$, or three-tenths of one percent.

Household sector and nonprofit organizations

In both the Flow of Funds (FOF) accounts produced by the Federal Reserve Board (FRB) and the National Income and Product Accounts (NIPA) produced by the Department of Commerce, activities undertaken by non-profit organizations are included in the household sector, as nonprofits are assumed to act on behalf of households.

Net saving and holding gains data in the Flow of Funds accounts

The “Net Saving” data in table 1 comes from the “Net investment” line (line 2) of Table R.100 in the Flow of Funds report. The “Holding gains” data in table 1 equals the sum of “Holding gains on assets stated at market value” and “Holding gains on assets stated at current cost” (lines 9 and 16, respectively) of Table R.100.

Net saving and purchases of consumer durable goods

The Flow of Funds treats consumer purchases of durable goods as a type of household

⁵³ Households are net purchasers of corporate equities through mutual funds and various retirement accounts; but this has been offset primarily by equity retirements associated with mergers and acquisitions as well as sales of equities to the foreign sector.

investment; net investment for this category of tangible assets is measured by expenditures of households on durable goods (gross investment from line 20 of table F.9 of the Flow of Funds report) minus estimated depreciation on existing durable goods (line 25 of table F.9). In the NIPAs, purchases of consumer durable goods are treated purely as expenditures in constructing the Commerce Department's estimate of personal saving. Recently, NIPA personal saving (shown in the memo item on line 6 of table 1 in the paper) has been smaller than "Net saving" in the Flow of Funds (line 3 of table 1).

Stock market wealth: our definition and the Flow of Funds definition

Line 6 of table B.100.e of the Flow of Funds accounts, "Balance Sheet of Households and Nonprofit Organizations with Equity Detail (1)," lists household holdings of corporate equities stated at market value. We define stock market wealth held by households slightly differently: Our definition adds the market value of corporate equities in private defined-benefit pension plans and equities in pension plans of state and local government employees to the Flow of Funds value.

Specifically, we define stock market wealth as the sum of: (a) household's holdings of corporate equities (B.100 line 24); (b) corporate equities held by private pension funds (L.119 line 14); (c) corporate equities held by state and local government retirement funds (L.120 line 13); (d) corporate equities held by bank personal trusts and estates (L.116 line 14); (e) corporate equities held by closed-end funds (L.123 line 6); (f) corporate equities held by mutual funds (L.122 line 10); and (g) corporate equities held by life-insurance companies (L.117 line 13), multiplied by the ratio of reserves of life insurance companies (L.117 lines 18 and 19) to the total financial assets of life insurance companies (L. 117 line 1). We define non-stock market wealth as total household net worth estimated by the Flow of Funds (B.100, line 43) minus our definition of stock market wealth. All tables come from the Flow of Funds.

Note that strictly adopting the Flow of Funds stock market series produces very similar results to those reported in the text.

II. The Life Cycle Model of Household Spending

Borrowing constraints

If households are unable to borrow against expected future income receipts to finance their desired current consumption plans, then the time path of spending might be more closely linked to the time path of income than in the simple life cycle theory. For example, in contrast to the strict life cycle implication, consumer spending might move with increases in income and wealth, even if the increases had been readily predicted by households ahead of time. In both household-level and aggregate data, it appears that consumption moves with predictable changes in income (Carroll and Summers (1991) provide a thorough analysis). This type of finding has led many economists to "reject" the plausibility of the most restrictive versions of the life cycle model, and to include forecasts of income growth in regression equations for consumption growth of the type

we present in section IV.⁵⁴

The other theoretical extensions to the basic life cycle model have been treated in detail by Muellbauer and Lattimore (1995).

Deriving the propensity to consume

In our simple examples, the formula given by equation (2) with $m_t = (1/\text{number of remaining periods in the lifetime})$ is a *solution* to the benchmark life cycle spending problem. This means that plugging in data for H_t (if they were available) and W_t on the right-hand side of equation (2), then multiplying by the propensity to consume m_t delivers the optimum choice for consumption.

We achieve this particular solution because interest is not paid on assets and people value their future consumption, today, as much as they value today's consumption. Muellbauer (1994) derives the analogous mathematical result in a more general life cycle framework that includes interest payments, inflation, and the "discounting" of future consumption in considering today's choice. In the more general framework, equation (2) still describes the optimum consumption choice: spending equals the propensity to consume multiplied by the sum of human and financial wealth. However, the exact formula for m_t includes rates of interest, inflation, and discounting (along with the number of remaining periods in the lifetime).

Realistic versions of the life cycle model

The basic results and insights from example 1 – the simple life cycle model – survive if we introduce time-varying consumption "needs" into the model; these could capture costs associated with raising young children or financing college tuition, for example. Also, we could assume that the person receives a moderate pension in old-age to help finance consumption without qualitatively affecting the analysis. Finally, introducing inheritances and bequests might not affect the model's results very much, especially if bequests are considered to be a close substitute for an old person's own consumption.

⁵⁴ Aside: This finding also has generated a subsequent debate among economists regarding whether borrowing constraints or precautionary saving motives explain why consumption "tracks" income in the data. Browning and Lusardi (1996) survey the diverse arguments.

III. Estimating Wealth Effects

Accounting for changes in the consumption-income ratio

Let g_c and g_y denote the growth rates of consumption and income, respectively, implying $C_t = (1 + g_c)C_{t-1}$ and $Y_t = (1 + g_y)Y_{t-1}$. The change in the consumption-income ratio can be decomposed as follows:

$$\begin{aligned}\Delta\left(\frac{C_t}{Y_t}\right) &= \frac{C_t}{Y_t} - \frac{C_{t-1}}{Y_{t-1}} = \frac{C_{t-1}(1 + g_c)}{Y_t} - \frac{C_{t-1}(1 + g_y)}{Y_{t-1}(1 + g_y)} \\ &= \frac{C_{t-1}(1 + g_c) - C_{t-1}(1 + g_y)}{Y_t} \\ &= \frac{C_{t-1}}{Y_t}(g_c - g_y)\end{aligned}$$

Cointegration of consumption, income and wealth data

A formal (and technically precise) discussion of cointegration is beyond the scope of this paper, but we can provide some introduction to the issue. See Hamilton (1994) for the details. Speaking intuitively, by themselves the macroeconomic series C_t , Y_t , and W_t display substantial upward – but unpredictable – trends over time; each of the series, therefore, is said to be *nonstationary*. In general, regression coefficients cannot be estimated by applying ordinary least squares to equations with nonstationary variables because the regression errors will themselves tend to be nonstationary. However, cointegrated series have the property that despite each variable being nonstationary, a regression of one variable on the others produces a stationary error, i.e., an error that tends to stay relatively small *and* reverts toward zero over time. Specifically, for a set of variables to be cointegrated, each variable in the system must be non-stationary and a linear combination of the variables must produce a stationary error. Cointegration is particularly desirable for applied macroeconomic research because, among other properties, estimates of the coefficients of an equation with cointegrated variables converge to the true parameter values faster than estimates from models with variables that are not cointegrated.

In the paper, we present results for several versions of equations (8) and (9). All of the coefficients reported come from equations which satisfy – statistically speaking – the conditions for cointegration: Nonstationary macroeconomic series, which in linear combination produce stationary regression errors.

We present coefficient estimates based on the Stock-Watson estimation procedure. A number of different valid statistical techniques exist to estimate the coefficients of an equation with cointegrated variables. One common alternative, the Johansen procedure,

usually produces parameter estimates quite similar to the Stock-Watson procedure; however, our experience indicates that the Johansen procedure is sensitive to somewhat arbitrary specification choices. For details on these procedures (and on cointegration more generally) see Hamilton (1994), chapters 19 and 20.

Interpretation of the coefficient on wealth in models 1 and 2

In Model 1, the coefficient b measures the response on the level of consumption, notated ΔC , to a one dollar change in wealth, ΔW . In contrast, in model 2, β measures the change of the logarithm of consumption, $\Delta \log(C)$, to a one unit change in the logarithm of wealth, $\Delta \log(W)$. One mathematical property of $\Delta \log(X)$ is that it approximately equals $\Delta X/X$. Therefore, the coefficient β of model 2 captures the response of $\Delta C/C$ to $\Delta W/W$. Since $\Delta C = (b \cdot \Delta W)$ from model 1 and $\Delta C = C \cdot (\beta \cdot \Delta W/W)$ from model 2, the only way for ΔC to be the same in both models is for $(b/\beta) = (C/W)$.

Log-difference of consumption as an approximation to the growth rate

If the level of consumption, C , grows at rate g between periods t and $t+1$, then $C_{t+1} = C_t(1+g)$. Dividing both sides by C_t and then taking natural logarithms gives $\log(C_{t+1}/C_t) = \log(C_{t+1}) - \log(C_t) = \Delta \log(C_{t+1}) = \log(1+g)$. Finally, for small g , $\log(1+g) \approx g$.

Predicted value of the level of consumption from the logarithms model (model 2)

In model 2 (the logarithms model), the target life cycle level of consumption equals $\exp\{\log(\hat{C}_t) + 0.5\sigma^2\}$, where σ is the estimated standard deviation of ϵ_t^2 .

Proportionality of human wealth to current income: household and aggregate data

Given the tendency for wages to rise with experience in the labor market and then to fall off sharply upon retirement, the assumption of a constant factor of proportionality between current and expected future incomes is not valid for an individual. However, if the age-distribution of aggregate wealth in the economy is stable (as we are assuming), then aggregate expected future income should be nearly proportional to current aggregate income. One consequence of this restriction is that the assumption of proportionality does not allow households to accumulate wealth to offset downward revisions to future expected income, as strict life cycle behavior requires.

Stock-Watson procedure for estimating cointegration coefficients

The Stock-Watson procedure involves adding leads and lags of the growth rates of each of the cointegrated series (the growth rates are themselves stationary) to equation (8) and (9), then applying ordinary least squares to the augmented regression.

IV. Sensitivity of Estimated Wealth Effects

Consumption errors from the six variants of the life cycle equations

Figure 11 plots the consumption errors (in percentage terms) arising from the three pairs of variants of the life cycle consumption equations – (1) and (2), (1a) and (2a), and (1b) and (2b). Comparing the two series *within* each panel illustrates that the “levels” and “logs” versions of the life cycle models generate nearly identical fits to the consumption data over history, despite generating different propensities to consume out of income and wealth. Comparing *across* the panels in Figure 11 illustrates that the different pairs of life cycle model variants produce fairly different time series of consumption errors.

Factors affecting estimated error correction speeds for consumption

Wealth information used to construct the life cycle error (ϵ_{t-1}). To predict the consumption growth rate in 2000:Q1 (for example), Brayton and Tinsley (1996) (BT) measure the life cycle error using net worth valued at December 31, 1999; Ludvigson and Steindel (1999) (LS) use net worth valued at September 30, 1999. Meanwhile, both BT and LS evaluate the life cycle errors using consumption and income data measured as in the national accounts (on a quarterly average basis). Our results in tables 6, 7, 8, and 10 are based on measuring all three series on a quarterly average basis: That is, consumption and income are measured as in the national accounts, and we take a geometric average of household net worth measured at the end of consecutive quarters.

In practice, we have found larger and more statistically significant error correction speeds for consumption when the life cycle error is constructed using the most recent data on household net worth; as well, using the most up-to-date wealth information tends to improve the fit of the short run equations for consumption growth. Our finding of sensitivity along this dimension is somewhat surprising, in the sense that graphs of life cycle errors constructed using net worth valued at the end of the previous quarter look very similar to those using net worth valued two quarters earlier. The correlation between the two quarterly life cycle errors is typically around 0.88, which apparently makes them different enough to significantly affect the estimated error correction speeds.

Note that the timing of the measurement of net worth does not appreciably affect estimates of the coefficients in the long run equations; hence, this judgment is not important for quantifying long run wealth effects.

Information used to measure the consumption response to income growth. Beginning with Flavin (1981), researchers have noted a tendency for growth rates of consumption and income to be more closely linked over the short run than is consistent with strict life cycle spending behavior.⁵⁵ As a consequence, it is common for researchers to include a

⁵⁵ Recall, strict life cycle behavior implies that consumption growth should respond only to unpredictable growth in income; researchers typically report, however, a tighter link between movements in spending and income – whether predictable or not – than is consistent with the simple theory. We

variable in equation (11) that captures the *excess sensitivity* of consumption growth to realized income growth.

As mentioned, using predicted income growth from a “fully specified” econometric model to capture the excess sensitivity of consumption yields substantially larger error correction speeds than using only last quarter’s measured income growth. Including the more accurate measure of current income growth in the equation for consumption growth transforms it closer to the spending-income ratio form of the short run equation, such as that shown in the upper row of table 8. To see this, suppose for simplicity that the short run consumption growth equation is of the following form

$$\Delta \log(C_t) = \gamma_1 + \gamma_2 \epsilon_{t-1} + \gamma_3 \Delta \log(\hat{Y}_t).$$


Rearranging variables produces

$$\Delta \log(C_t) - \gamma_3 \Delta \log(\hat{Y}_t) = \gamma_1 + \gamma_2 \epsilon_{t-1}.$$

Note that the left hand side of this equation equals the difference of the growth rate of consumption and the weighted growth rate of income, denoted $g_c - \gamma_3 g_y$. When γ_3 is close to zero (as it is when we use the lagged growth rate of income as our $\Delta \log(\hat{Y}_t)$ variable), then the left hand side simply collapses to g_c . However, if γ_3 is close to 1.0, as it is when we use the full econometric model to produce $\Delta \log(\hat{Y}_t)$, then $\Delta \log(C_t) - \gamma_3 \Delta \log(\hat{Y}_t)$ approximately equals $g_c - g_y$. As shown in the technical appendix item *Accounting for changes in the consumption-income ratio*, $\Delta(C_t/Y_t)$ also approximately equals $g_c - g_y$. We have already demonstrated the tendency for larger error correction speeds in the spending-income ratio form compared with the consumption growth form: see the top row of table 8. Therefore, it is not a surprise that estimated error correction speeds rise when we use the full econometric model for $\Delta \log(\hat{Y}_t)$.


discuss the excess sensitivity of consumption and income – and provide some research citations – in the technical appendix item *Borrowing constraints*.

FIGURE 1
Wealth-Income Ratio^a and Personal Saving Rate
1953:Q1 - 2000:Q1




a. Ratio of net worth of households, as measured in the Federal Reserve Board's Flow of Funds Accounts, to disposable personal income, as published by the Bureau of Economic Analysis.

FIGURE 2
Households' Real Rate of Capital Appreciation on Non-Stock Market Wealth^b
1953 - 1999




b. The dashed line occurs at 0.43%, the average real rate of appreciation of non-stock-market wealth from 1953 through 1999.

FIGURE 3
Households' Real Rate of Capital Appreciation on Stock Market Wealth^c
1953 - 1999




c. The dashed line occurs at 7.75%, the average real rate of appreciation of stock market wealth from 1948 through 1999.

FIGURE 4
Households' 5-Year Cumulative Real Rate of Capital Appreciation
on Non-Stock Market Wealth^d
1958 - 1999



d. Refers to the real 5-year cumulative rate of appreciation *ending* at the listed year.

FIGURE 5
Households' 5-Year Cumulative Real Rate of Capital Appreciation
on Stock Market Wealth^e
1953 - 1999



e. Refers to the real 5-year cumulative rate of appreciation *ending* at the listed year.

FIGURE 6
Actual Consumption and Consumption Predicted from Models 1 and 2
1960:Q1 - 2000:Q1




FIGURE 7
Actual Consumption and Consumption Predicted from Models 1 and 2
1995:Q1 - 2000:Q1




FIGURE 8
Consumption Errors Predicted from Models 1 and 2
(1960:Q1 - 2000:Q1; measured in percent)





FIGURE 9

Wealth Effects under Alternative Adjustment Rates for Actual to Planned Consumption



Alternative consumption paths:

- C(I): Short-run wealth effect shown in table 7
- C(II): Consumption fully responds to wealth within the first period
- C(III): Same immediate response as C(I); faster error correction thereafter
- C(IV): Same immediate response as C(I); no error correction

Figure 10
Consumption Errors from Selected Variants of Life Cycle Equations
(1994:1 to 2000:1; measured in percent)




FIGURE 11

Consumption Errors from the Six Variants of the Life Cycle Equations
(1960:Q1 - 2000:Q1; measured in percent)

