The Consumption Wealth Effect

Charles Steindel[1]

Federal Reserve Bank of New York

The last decade has seen extraordinary swings in the aggregate value of the net worth of American households. According to the Federal Reserve's Flow of Funds statistics, aggregate household wealth rose nearly $18 trillion, or almost 75%, from the end of 1994 to the end of 1999. Wealth then fell about 7%--about $3 ½ trillion—through the end of 2002. Since then, wealth has grown more than 15%, or over $7 trillion, to again reach a new high well above $47 trillion.[2]

The personal saving rate is now virtually zero, whereas it was above 7% in the early 1990s (Chart 1). Taking this longer perspective, the pronounced increase in wealth relative to income since the mid-1990s appears to be linked to the reduction in the saving rate. This broad correlation has often been cited as confirmation of a causal link between increases in wealth and increases in consumer spending (other things equal, if consumer spending rises, the saving rate falls).

As will be shown, the increase in wealth has primarily reflected sharp gains in the price of equities and, in recent years (to a lesser extent than is commonly understood), housing. There are concerns that if these prices were to decline consumers would then reduce spending substantially in order to "restore" their balance sheet positions.

The focus on the role of wealth movements in determining spending stems directly from the original work on the life cycle model of consumer spending of Franco Modigliani and his collaborators (for instance, Ando and Modigliani, 1963; Modigliani, 1971). In the typical empirical formulation of the model, consumer spending has been modeled as a linear function of income and wealth; with the coefficient on wealth typically found to be in the neighborhood of .05. Some formulations (for instance, Modigliani, 1971; Modigliani and Steindel, 1977) have allowed for a lagged response of consumer spending to wealth changes. In any event, wealth changes in the multiple trillions of dollars would ultimately generate consumer spending changes in the hundreds of billions of dollars. The annual peak level of personal saving was the $366 billion recorded in 1992, so it seems safe to assert that wealth increases have been the primary driver of the decline in saving since the mid-1990s and likely a major

impetus behind consumer spending in this period. Correspondingly, fears of the consequences of a retreat in wealth look to be well-founded.

Recent research (Ludvigson and Steindel, 1999; Lettau, Ludvigson, and Steindel, 2002; Lettau and Ludvigson, 2001, 2004) suggests that this line of reasoning is a bit oversimplified. In essence, the long-run correlations between wealth, consumption, and saving suggested by the life cycle hypothesis clearly exist. However, it is incorrect to apply these long-run relationships to short-term movements; pronounced changes in wealth need not be associated with any substantive change in spending or saving. It is even inappropriate to assume that long-sustained changes in wealth can always be fully identified with the long-run changes associated with major revisions to spending plans. There is indeed a connection between the increase in wealth over the last decade and the decline in saving. In all likelihood there have been forces (such as the higher productivity trend) that have worked both to increase market valuations and boosted consumer estimates of permanent income relative to realized income, thus increasing spending and reducing saving. As the failure of consumer spending to retreat during the 2000-2002 bear market suggests, these fundamental reassessments of the longer-term outlook can be largely independent of even very large moves in current market valuations.

The next section discusses a bit further the concept of the consumption-wealth effect, noting the somewhat different ways the term is used. This is followed by a discussion of the uses of wealth changes as a forecasting instrument for consumer spending and a further explanation of the recent line of research on the short- and long-term connections between consumer spending, income, and wealth.

What is the Consumption-Wealth Effect?

Real aggregate consumer spending and wealth are clearly correlated over long horizons (Chart 2). What is apparently meant, though, in most discussions of the consumption-wealth effect, is not this fairly obvious linkage, but one of two somewhat related shorter-term connections:

1. Movements in wealth have a predictable short-term effect on movements in spending.
2. Movements in wealth can be used to predict short-term movements in spending.

Strictly speaking, these two usages can be independent. A structural connection between wealth changes and spending may exist, but not be useful for forecasting consumption, simply because too many

other factors jointly affect both. Conversely, there could be factors influencing wealth and, with a lag, consumption, leading to a nonstructural predictive relationship between the two.

These nonidentical, but clearly related, uses of the term consumption-wealth effect should be distinguished from different usages seen in the literature. Some studies (for instance, Poterba and Samwick, 1995; Starr-McCluer; 1998; Poterba, 2000; Dynan and Maki, 2001) focus on whether particular changes in wealth (and often on changes in particular types of wealth) were associated with changes in spending (and in some cases changes in particular types of spending. For instance, Poterba and Samwick, 1995, and Poterba, 2000 looked at changes in spending on luxury goods in the wake of stock market moves). This is different from attempts to measure the (structural or forecasting) linkage between wealth changes and aggregate spending. Finding that people who have become richer will subsequently increase spending, or calibrating that relationship for particular episodes, is not the same as finding that there is a stable relationship in the aggregate data between increases in wealth and increases in spending. This paper focuses on the evidence for the latter relationship; the "consumption-wealth effect" here means a well-measured stable connection between aggregate wealth movements and aggregate consumption. Typically a major interest is the ability to exploit such a relation for macroeconomic forecasts, but also there is an interest in gauging the size of any structural behavioral connection that may exist, for purposes such as understanding potential trends in capital formation.

The Sources of the Wealth Effect

The linkage between wealth and consumption is most explicitly examined in the context of the life cycle model of the consumer spending, developed by Franco Modigliani and collaborators (Major works in its development include Brumberg and Modigliani, 1954, 1978; Ando and Modigliani, 1963; Modigliani, 1966). In essence, the life cycle model starts from premise of the permanent income model that consumers will keep spending smooth in the face of changes to income. The model is enriched by positing that the major movement in income is the loss of wages with retirement. The anticipation of retirement implies that households will be accumulating wealth over the course of their working lives and dissipating that wealth in retirement (given the occasionally sharply questioned—e.g. Kotlikoff and Summers, 1981; Modigliani, 1988; Kotlikoff, 1988—assumption that estate planning is fundamentally different than other types of saving motives). Unless the rate of time preference is higher than the rate of return, households will also

plan for an increasing stream of spending. Thus, the basic logic of the life-cycle model is that there will be a positive aggregate association between wealth and consumer spending (at least in a nation with a growing population). Older households will own the bulk of a nation's wealth and do the lion's share of spending.

What will happen to a household when there is some change in wealth not connected to the spending plan (say by winning a lottery)? The logic of the model is that the long-run spending plan will be revised; for an increase in wealth there will be an upward adjustment in spending. Since the plan is assumed to involve spending all the increment to wealth over the remaining life expectancy, the increment to spending will be somewhat higher than the rate of return. Thus the logic of the usual "5 cents on the dollar" estimate of the marginal propensity to consume from an increment to wealth stems from the comparison of that figure with usual estimates of the long-run real risk-free rate of return (implicitly assumed to be the rate the invested windfall earns). Specifically, the consumer spending plan can be formulated as (assuming a zero rate of time preference)

(1) $C_t=(1-(1+r)^{-\sigma/(\sigma+1)})/(1-(1+r)^{-T\sigma/(\sigma+1)})*PV_t=\delta(r,T,\sigma)*PV_t$

Where

C_t=Consumer spending at time *t*.

σ= Elasticity of substitution between current and future consumption.

r= Real rate of interest

T= Life expectancy

PV_t= Resources available to the consumer at time *t*. These resources consist of the present value of income anticipated to be earned for the balance of the consumer's working life plus the current value of nonhuman wealth.

In general, if the elasticity of substitution is less than one, δ will be larger than $1/T$ and at low interest rates will be, very roughly, comparable to the rate of return.[3]

The seminal Ando and Modigliani (1963) paper estimated a model relating the annual U.S. data on consumer spending to income and wealth and reported a coefficient on wealth of around .05—very much in line with the logic argued above. Some years later, Modigliani (1971) extended the specification to quarterly data and found similar results. Most importantly, Modigliani found that to a first approximation consumer spending responses to changes in wealth were largely independent of the type of wealth that

changed—consumers respond to changes in both the stock market and in other forms of wealth. Also, Modigliani found that there was a lagged response of spending to wealth changes. This lag has major consequences for economic forecasting. As will be noted below, the major element behind short-run changes in wealth in the United States is movements in the stock market. Stock market moves are continuously observed, and their consequences for household wealth are readily computed. Thus the existence of a stable lagged effect from wealth to consumption provides a major mechanism to forecast consumer spending and thus economic activity as a whole. Equations of the type Modigliani developed remain important components of widely-used commercial econometric models.

Criticisms and Modifications of the Wealth Effect

The macroeconomic wealth effect as outlined in the literature on the life cycle model would appear to be a powerful combination of well-reasoned theory backed by compelling statistical evidence. Nonetheless, it has been subject to a number of criticisms, and some recent literature suggests that major modifications are in order. These modifications may be of sufficient magnitude to largely nullify the value of the model as a tool of macroeconomic forecasts, even if one accepts its basic premises. Before considering these modifications, some more traditional lines of criticism will be addressed.

Criticism 1: Simultaneity

The canonical estimates of the aggregate life-cycle equation (for instance, those in Ando and Modigliani, 1963, and Modigliani, 1971) were made using single equation techniques (ordinary least squares, or ordinary least squares modified to correct for serially correlated errors). There is clear reason to suspect that consumption innovations have a direct effect on wealth, thus raising the possibility that the wealth coefficient is biased.[4]

This criticism is obviously valid, but its force (at least in terms of bias as usually gauged in linear systems) may be modest. The connection from consumption innovations to wealth is likely positive. There is likely a similar positive connection from consumption innovations to income, the other major determinant of spending in these models. Given these common linkages, the amount of bias in the estimated coefficients on income and wealth in a linear regression may be inconsiderable. Indeed, limited estimates of such models (e.g., Mishkin, 1976, for a model of consumer durable spending) using simultaneous equation techniques found quite similar results to those using single equation methods.

Criticism 2: Wealth is too narrowly Held

The bulk of U.S. wealth is owned by a small part of the population (Tracy, Schneider, and Chan, 1999, Poterba, 2000, and Maki and Palumbo, 2001, show some of the recent data wealth distribution by type.). A natural criticism of the life cycle model is that wealth changes directly impacting so few people can not reasonably be expected to affect spending by the population at large. The observed connection between wealth changes and spending may be an artifact resulting, perhaps, from swings in consumer confidence that also has ramifications for the stock market.

Fundamentally this latter argument is a focusing of the simultaneity criticism. It looks flawed on a number of grounds. While clearly "consumer confidence" must be a determinant of spending, there is at most rather limited evidence that observed consumer confidence measures have any explanatory power for spending over and above income and wealth (Modigliani 1971, Mishkin, 1977, Bram and Ludvigson, 1998, Ludvigson, 2004). More importantly, there is confusion in this argument between changes in spending by large numbers of people and large changes in spending. The more affluent segment of the population, which holds the bulk of the wealth, also accounts for a disproportionate amount of spending.

Criticism 3: Wealth Effects should be Disaggregated (or, Housing is Different)

Standard estimates of life cycle models either aggregate all wealth as one entity or, at most, following the example of Modigliani, 1971, separate wealth into stock market and nonstock components. This looks troubling on a number of grounds:

a. The pure life cycle wealth effect is sensitive to the age of the consumer experiencing a change in wealth (an older consumer should change spending more). If there is a systematic demographic difference in the distribution of the consumer portfolio across age groups by type of assets the assumption of a common fixed effect looks to be incorrect. Even looking past demographics, changes in the value of housing could conceivably have radically different effects on spending than changes in other forms of wealth. One argument is that because housing is much more evenly held across the population than other types of wealth (Tracy, Schneider, and Chan, 1999) there may be a stronger connection between changes in home values and changes in spending. More credit-constrained consumers own homes than own other assets, thus a rise in home values eases credit constraints for a large number of households, further spurring spending. This

argument is reinforced by the increased ease of refinancing home equity; the ability to tap the wealth from homes may have increased the propensity to consume from changes in housing values, reducing easing credit constraints for a growing number of homeowners (Canner, Dynan, and Passmore, 2002, discuss possible effects of increased refinancing on household behavior).

b. In striking contrast is an argument that increases in home values should have very limited effects on spending. In this line of reasoning, increases in existing home values are not an increase in aggregate wealth, since they just raise the cost of acquisition to non-owners, who then suffer a capital loss (unmeasured) equal to the gains made by the owners.

Before addressing these criticisms some observations can be made about the dynamics of aggregate wealth in the United States. Table 1 shows the distribution of household wealth by type at the end of 1973, 1982, 1994, 1999, 2002, and 2003. As a rough approximation, aggregate equity exposure (including that held by mutual funds and fiduciaries) has often been equal to or larger than aggregate home values, but clearly equity has fluctuated more. Chart 3 shows how capital gains on equity typically dominate changes in aggregate household wealth, even over longer periods, when cumulated saving flows would be expected to play a more important role. It simply is not the case that real estate fluctuations typically determine aggregate wealth movements. Considering the dynamics of wealth movements alone some case might be made for distinguishing between stock market and nonstock market wealth effects, but signaling out housing does not look so clear-cut.[5]

Looking more closely at the experience of the last decade, the major increase in wealth in the latter of the 1990s was clearly due to the bull market in stocks, and the shrinkage in wealth in the early part of this decade reflected the retreat in the market. The sharp rise in housing values only stemmed the loss in wealth; they did not eliminate it. The more recent turnaround in wealth reflected the recovery of the stock market being added to the housing gains—though it is clear that the new peak in wealth has occurred even though the stock market remains below its peak, and would not have occurred without the gains in housing. In other words, the period since 2000 is a bit unusual in that since that there has been, on balance, further increases in real wealth while the stock market has been down—but this is reflects a very unusual choice of dates. The general rule that the stock market is the major force behind wealth movements looks intact.

Returning the housing-related critiques of the wealth effect evidence, the first one (the distribution of wealth should matter) is surely sound. However, as is usual with criticisms on disaggregated grounds of aggregate results it is a bit hard to address. Reliable time series data on the distribution of wealth holdings are simply not available in a timely fashion, so as a practical matter it looks a bit extreme to object to results, appearing reliable on other grounds, solely on the basis of what may be a weak justification for aggregation. The argument that the growth of the refinance and home equity market has systematically increased the propensity to consume from housing is more complex. In principle, there could well be a connection between the liquidity of an asset and its relation to spending, but the sign is ambiguous on usual income and substitution effect grounds. There could well be a linkage between use of home equity borrowings and other mortgage financings and spending (Canner, Dynan, and Passmore, 2002, find that an appreciable fraction of funds raised by these methods do go to finance consumer spending), but this may be due to the growing convenience of this method of finance, rather than a stronger behavioral link from changes in home values to spending.

The gist of the second criticism is that there can't be a "wealth effect" from housing because gains to homeowners are implicitly matched by increases in the user cost of housing. The present value of the increased user cost will exactly match the increase in the value of housing; therefore there is no expansion of the aggregate budget line.[6] The flaw in this argument is the unstated assumption that the technical ability of the housing stock to provide shelter is fixed. In principle, households may offset some of the increased user cost by explicitly renting out part of their housing stock, leaving a net gain in explicit income. While households may not in practice do this the opportunity is there and thus the ability to realize the aggregate gain.[7]

On balance, there looks to be rather little reason to anticipate a radically different link from housing to spending than for other components of the balance sheet, and life cycle results presupposing a uniform response (or not separating housing) should not be rejected on these grounds alone.

A Modified Wealth Effect

While the more "traditional" criticisms of the aggregate life cycle wealth effect look to have little net force, more recent work suggests that the model needs to be modified. Dissatisfaction with the standard

version of the model stems from experiences of the 1990s and modern thought on the integration of asset market valuation and forward-looking consumer behavior.

The latter part of the 1990s did see enormous increases in household wealth and a drop in the personal saving rate, as well as some acceleration in the growth of consumer spending. These are all broadly consistent with the life cycle model. However, the magnitude of the change in consumer spending in response to the change in wealth was substantially smaller than would likely have been predicted. Between year-end 1994 and 1999, "real" capital gains (gains over and above those needed to maintain the purchasing power of wealth[8]) cumulated to about $13 trillion. Assuming a .03 propensity to consume from such gains, consumption should have risen roughly $400 billion relative to measured disposable income over this period; this would have corresponded to a drop in the personal saving rate of at least 6 percentage points. Instead, the actual decline in the saving rate was in the vicinity of 2 percentage points. The propensity to consume from the increase in wealth was apparently very low. An ad hoc rationale for this low propensity is that there is some floor on the saving rate. It is not clear why such a floor would exist (as a matter of logic gross investment in a closed economy can not fall below zero, but gross investment is a very different concept than the U.S. personal saving measure). However, even if we accept the existence of such a floor, then the 2000-2003 experience also suggests very low propensities to consume from wealth. "Real" capital losses totaled around $6 trillion over those three years; with a .03 propensity to consume from wealth consumption growth should have lagged income growth by around $200 billion, pushing up the personal saving rate several points. Instead the saving rate was flat to down over these years.

Clearly, analysis of household spending and saving behavior over the last decade based on observation of aggregate wealth changes and the application of "traditional" propensities to consume from wealth movements need to confront the reality that consumption and saving have been more stable than this model would suggest. Alternatively, it looks as though the propensity to consume from wealth was low in recent years. Chart 4 illustrates movements in the ex-post propensity to consume from wealth. The top figure illustrates the propensity estimated from rolling 20-year regressions relating real per-capita consumer spending to disposable income and wealth; the other two figures come from similar rolling regressions in which wealth is decomposed into stock and non-stock components.[9] There is striking variation in the

propensities to consume from wealth, with the estimates for the total and for equity over recent periods being very low indeed. If one wishes to use observed changes in wealth (or those assumed to occur in the future) to project consumer spending it is not at all clear what value to choose for this propensity. In an era when annual changes in wealth can easily amount to several trillion dollars, small changes in the assumed wealth propensity can have major implications for a projection of consumption.

In sum, the experiences of the last 10 years may illustrate some substantive shortcomings of the life cycle model's traditional empirical aggregate framework (not the logical underpinnings of the model) and the use of a fixed wealth effect to explain and forecast consumer spending. Modifications begin by re-examination of the basics of the model Consumers are assumed to adjust their spending path in reaction to a change in wealth. However, an unstated assumption is that the change in wealth has no possibility of being reversed. This makes sense if the wealth change is the result of a lottery win. The situation is very much different in the case of swings in the aggregate value of the stock market, which history suggests have a strong likelihood of being reversed within a consumer's planning horizon.[10] The logic of the forward-looking model is that consumers should in the aggregate respond only to swings in wealth values that are perceived to be "permanent."

This argument has been long-accepted, and to some the lag found in traditional life cycle models between aggregate wealth changes and changes in spending may be a sign that consumers are reacting only to "permanent" changes in wealth. However, estimating the permanent component of wealth changes through fixed lags is inadequate; the stock market, at least, appears to go through long periods of "over" and "under" valuation. A more coherent metric for determining the permanent component of wealth changes is desirable.

In a series of papers Sydney Ludvigson and collaborators (Ludvigson and Steindel, 1999; Ludvigson, Steindel, and Lettau, 2002; Lettau and Ludvigson, 2001, 2004) have developed techniques for estimating this "permanent" component of wealth. It can be shown that wealth, consumption, and income are cointegrated, precisely as the life cycle model predicts (and the coefficients of the cointegrating vector are consistent with the long-standing views of the model). The statistical evidence suggests that, in the vector-error correction formulation of this system, innovations of wealth do not precede innovations in either consumption or income. Thus, a natural interpretation is to view the innovations of wealth estimated

from the cointegrating vector as "transitory" changes in wealth, with little consequences for spending, while the balance of wealth changes (which will be associated with contemporaneous moves in spending and income) are "permanent." In this reading of the evidence, there is no "lagged" wealth effect; there is an immediate association of permanent wealth changes with changes in spending. The lagged effects seen between aggregate changes in wealth and aggregate changes in spending are an artifact of misspecification of the traditional equations.

In this reading of the data there was a decided "wealth effect" in the 1990s (since a cointegrated system is not precisely structural the classical interpretation is a bit problematic, but convenient); permanent wealth did increase sharply, with the rise possibly associated with the increase in productivity growth (thus likely boosting permanent labor income as well), and was associated with the increase in spending. However, much of the upward movement in wealth was not permanent, and was not perceived as such by consumers, and did not affect spending—even though the stock market stayed very high for a prolonged period. Thus, the wealth effect gauged by observation of aggregate wealth changes and movements in spending appeared to be low. When the bull market broke after 2000, consumers did not revise down their estimates of permanent wealth and maintained spending; conversely, the recovery of the market starting in 2003 also did not radically change long-term views. Chart 5 compares actual movements in wealth with estimates of the permanent component. It is clear that in the 1990s growth in permanent wealth lagged the total. Permanent wealth held steady after 2000 while overall wealth faltered for a number of years.

An intriguing aspect of this view of the wealth effect is that it stands the traditional timing on its head. Instead of observed movements in wealth leading movements in consumer spending, it appears that innovations in consumer spending lead changes in wealth (in other words, households may tend to increase spending when they adjust upward their longer-views on income and wealth dynamics). Increases in spending in a period when the market is low may be an indication that the market will improve. Aggregate consumption may be a useful forecasting device for the aggregate market.[11]

These results seem a bit startling, and at times have been viewed as dismissive of the existence of the "wealth effect." Actually, they are quite consistent with some very long recognized phenomena. Most notably, wealth is very much more volatile than income or spending (Chart 6). Given this, it is hard to

believe that the data can support a fixed high frequency linear relationship between the three. Also, the data come very close to supporting a random walk for consumption—it is nearly impossible to forecast changes in consumption growth. If consumption is highly difficult to forecast, why should one suppose that movements in wealth may help do so?

The cointegration-vector error correction model shows us that there is indeed a "wealth effect" in the permanent components of the series (not to be confused with low frequency moves in wealth). The wealth effect is in magnitude in line with that traditionally associated with the life cycle model. It is aggregate, not associated with any particular component of the balance sheet—but given the realities of asset market moves clearly the stock market plays a major role. The wealth effect is immediate, not lagging, and as such is not exploitable for unconditional macroeconomic forecasts (of course, conditional forecasts may be made assuming movements in the permanent component of wealth). Movements in wealth not associated with moves in consumption or income are transitory. They "ultimately" will not be sustained, but this timing is unpredictable.

Of course, this view of the wealth effect is not itself without its weaknesses. Statistically, it rests on the details of the cointegration-vector error correction relationship between wealth, income, and spending. In principle, further analysis of the data could overturn this result (for instance, Davis and Palumbo, 2001, argue that innovations in wealth precede innovations in consumption, while Rudd and Whelan, 2002, contend that cointegrating relationships can not be found for alternative concepts of wealth, income, and spending). The earlier criticisms of aggregated approaches would hold for this model as well; better disaggregated data and results could also overturn the result. But at the moment this modified view, which preserves the basic message of the life cycle model, but reduces the role of wealth movements in forecasting exercises, looks to be the most consistent with the evidence.

Table 1

Equity and Housing in the Household Balance Sheet

	Equity Share	Housing Share
1973	20.1%	27.2%
1982	13.4	30.9
1994	23.8	30.0
2000	41.8	24.2
2002	26.2	34.1
2003	30.5	33.5

Equity includes holdings of bank trusts, mutual and pension funds, and insurance companies.

Source: Federal Reserve Board Flow of Funds.

References

Ando, Albert, and Franco Modigliani (1963). "The Life Cycle Hypothesis of Saving: Aggregate Implications and Tests." American Economic Review, vol. 103, no. 1, pp. 55-84.

Brumberg, Richard, and Franco Modigliani (1954). "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in Post-Keynesian Economics, New Brunswick (Rutgers).

_____, and Franco Modigliani (1980). "Utility Analysis and Aggregate Consumption Functions: An Attempt at Integration," in The Collected Papers of Franco Modigliani: Vol. 2, The Life Cycle Hypothesis of Saving, Cambridge, Massachusetts (MIT).

Canner, Glenn, Karen Dynan, and Wayne Passmore (2002). "Mortgage Refinancing in 2001 and early 2002," Federal Reserve Bulletin, pp. 469-480.

Corrado, Carol, and Charles Steindel (1980). "Perspectives on Personal Saving," Federal Reserve Bulletin, vol. 66, no. 8, pp. 615-625.

Davis, Morris, and Michael Palumbo (2001). "A Primer on the Economics and Time Series Econometrics of Wealth Effects." Finance and Economics Discussion Series, no. 2001-9, Federal Reserve Board.

Dynan, Karen, and Dean Maki (2001). "Does Stock Market Wealth Matter for Consumption?" Finance and Economics Discussion Series, no. 2001-23, Federal Reserve Board.

Kotlikoff, Laurence (1988). "Intergenerational Transfers and Savings," Journal of Economic Perspectives, vol. 2, no. 2, pp. 41-58.

_____, and Lawrence Summers (1981). "The Role of Intergenerational Transfers in Aggregate Capital Accumulation," Journal of Political Economy, vol. 89, no. 4, pp. 706-732.

Lettau, Martin, and Sydney Ludvigson (2001). "Consumption, Aggregate Wealth and Expected Stock Returns," Journal of Finance, vol. 56, no. 3 pp. 815-49.

_____, and Sydney Ludvigson (2004). "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption," American Economic Review, vol. 94, no. 1, pp. 276-299.

Ludvigson, Sydney (2004). "Consumer Confidence and Consumer Spending," Journal of Economic Perspectives, vol. 18, no. 2, pp. 29-50.

_____, and Jason Bram (1998). "Does Consumer Confidence Forecast Consumer Expenditure? A Sentiment Index Horse Race," Federal Reserve Bank of New York Economic Policy Review, vol. 4, no. 2, pp. 59-78.

_____, and Charles Steindel (1999). "How Important is the Stock Market Effect on Consumption?" Federal Reserve Bank of New York Economic Policy Review, vol. 5, no. 2, pp. 29-52.

_____, Charles Steindel, and Martin Lettau (2002). "Monetary Policy Transmission Through the Consumption-Wealth Channel," Federal Reserve Bank of New York Economic Policy Review, vol. 8, no. 1, pp. 117-134.

Maki, Dean and Michael Palumbo (2001). "Disentangling the Wealth Effect: A Cohort Analysis of Household Saving in the 1990s." Finance and Economics Discussion Series, no. 2001-21, Federal Reserve Board.

Mishkin, Frederic (1976). "Illiquidity, Consumer Durable Expenditure, and Monetary Policy," American Economic Review, vol. 66, No. 4, pp. 642-654.

_____, (1977). "What Depressed the Consumer? The Household Balance-Sheet and the 1973-75 Recession," Brookings Paper on Economic Activity, no. 1, pp. 123-164.

Modigliani, Franco (1966). "The Life Cycle Hypothesis of Saving, the Demand for Wealth, and the Supply of Capital," Social Research, vol. 33 (summer), pp. 160-217.

____, (1971). "Monetary Policy and Consumption," in Consumer Spending and Monetary Policy: The Linkages. Boston (Federal Reserve).

_____, (1988). "The Role of Intergenerational Transfers and Life Cycle Saving in the Accumulation of Wealth," Journal of Economic Perspectives, vol. 2, no. 2, pp. 15-40.

_____, and Charles Steindel (1977). "Is a Tax Rebate an Effective Tool for Stabilization Policy?" Brookings Papers on Economic Activity, no. 1, pp. 175-209.

Poterba, James (2000). "Stock Market Wealth and Consumption," Journal of Economic Perspectives, vol. 14, no. 2, pp. 99-119.

_____, and Andrew Samwick (1995). "Stock Market Ownership Patterns, Stock Market Fluctuations, and Consumption," Brookings Papers on Economic Activity, no. 2, pp. 295-357.

Rudd, Jeremy, and Karl Whelan (2002). "A Note on the Cointegration of Consumption, Income, and Wealth." Finance and Economics Discussion Series, no. 2002-53, Federal Reserve Board.

Starr-McCluer, Martha (1998). "Stock Market Wealth and Consumer Spending." Finance and Economics Discussion Series, no. 1998-20, Federal Reserve Board.

Steindel, Charles (1981). "The Determinants of Private Saving," in Public Policy and Capital Formation, Washington (Board of Governors of the Federal Reserve System).

Tracy, Joseph, Henry Schneider, and Sewin Chan (1999). "Are Stocks Overtaking Real Estate in Household Portfolios?" Federal Reserve Bank of New York Current Issues in Economics and Finance, vol. 5, no. 5.

[1] Thanks to Andrew Zuppann for assistance, and Martin Lettau and Sydney Ludvigson for their data on permanent wealth. The views expressed in this paper are those of the author, and are not necessarily those of the Federal Reserve Bank of New York or of the Federal Reserve System.

[2] The currently reported value for the end of the third quarter of 2004 is $46.7 trillion. The handsome gains in the stock market in Q4 would place the year-end value substantially above that value.

[3] Steindel (1981) derives the expression. Table 1 of that paper shows that at a 4% interest rate δ will be between .0333 and .0556 for a consumer whose elasticity of substitution is one and has a forty-year life expectancy.

[4] Typically lagged (start of period) wealth is used as a variable in estimated consumption functions. This practice does not negate the simultaneity criticism, since investors may well forecast innovations in spending in valuing assets.

[5] Those knowledgeable about the construction of the aggregate wealth data may object. For most assets, current-dollar aggregate wealth is computed from cumulated investment flows, with three major exceptions: equity and land (where current market valuations are used) and physical capital, where current costs of reproduction are used. Of course, the major physical assets owned by households and unincorporated business (the "household sector") are structures. Thus, as a matter of logic the only wealth types (with some very minor exceptions) for which capital gains and losses are recorded in the aggregate wealth data are equities and real estate. So a separation of equity and real estate (combined or themselves separated) from all other wealth components may be warranted. However, note that these data consideration apply only to current-dollar capital gains and losses; clearly components of the balance sheet whose current-dollar prices are fixed suffer real capital losses as consumption prices rise.

[6] A simplistic alternative rendering of this argument is that homeowners can't realize the gains on their houses, because they still "need a place to live." Presumably borrowing against an asset that has appreciated in value is not realizing a gain! Another rendering is that the increased cost of future home acquisition by renters balances out the gain to homeowners.

[7] Of course, there may be some extreme example where every bit of housing is owner-occupied by dynasties, eliminating the possibility of an explicit rental market margin to determine aggregate gains.

[8] Corrado and Steindel (1980) explain the derivation of "real" capital gains.

[9] The regressions include current values and 3 lags of each variable and correct for first-order serial correlation of the errors.

[10] Of course, an individual consumer can always lock in a gain by selling out or hedging. But both moves are costly and are, of course, unavailable in the aggregate.

[11] These are results stemming from ex post observations of an equilibrium system. They are forecasts assuming continued market equilibrium, not an arbitrage opportunity.

Chart 1: Wealth and Saving Rates

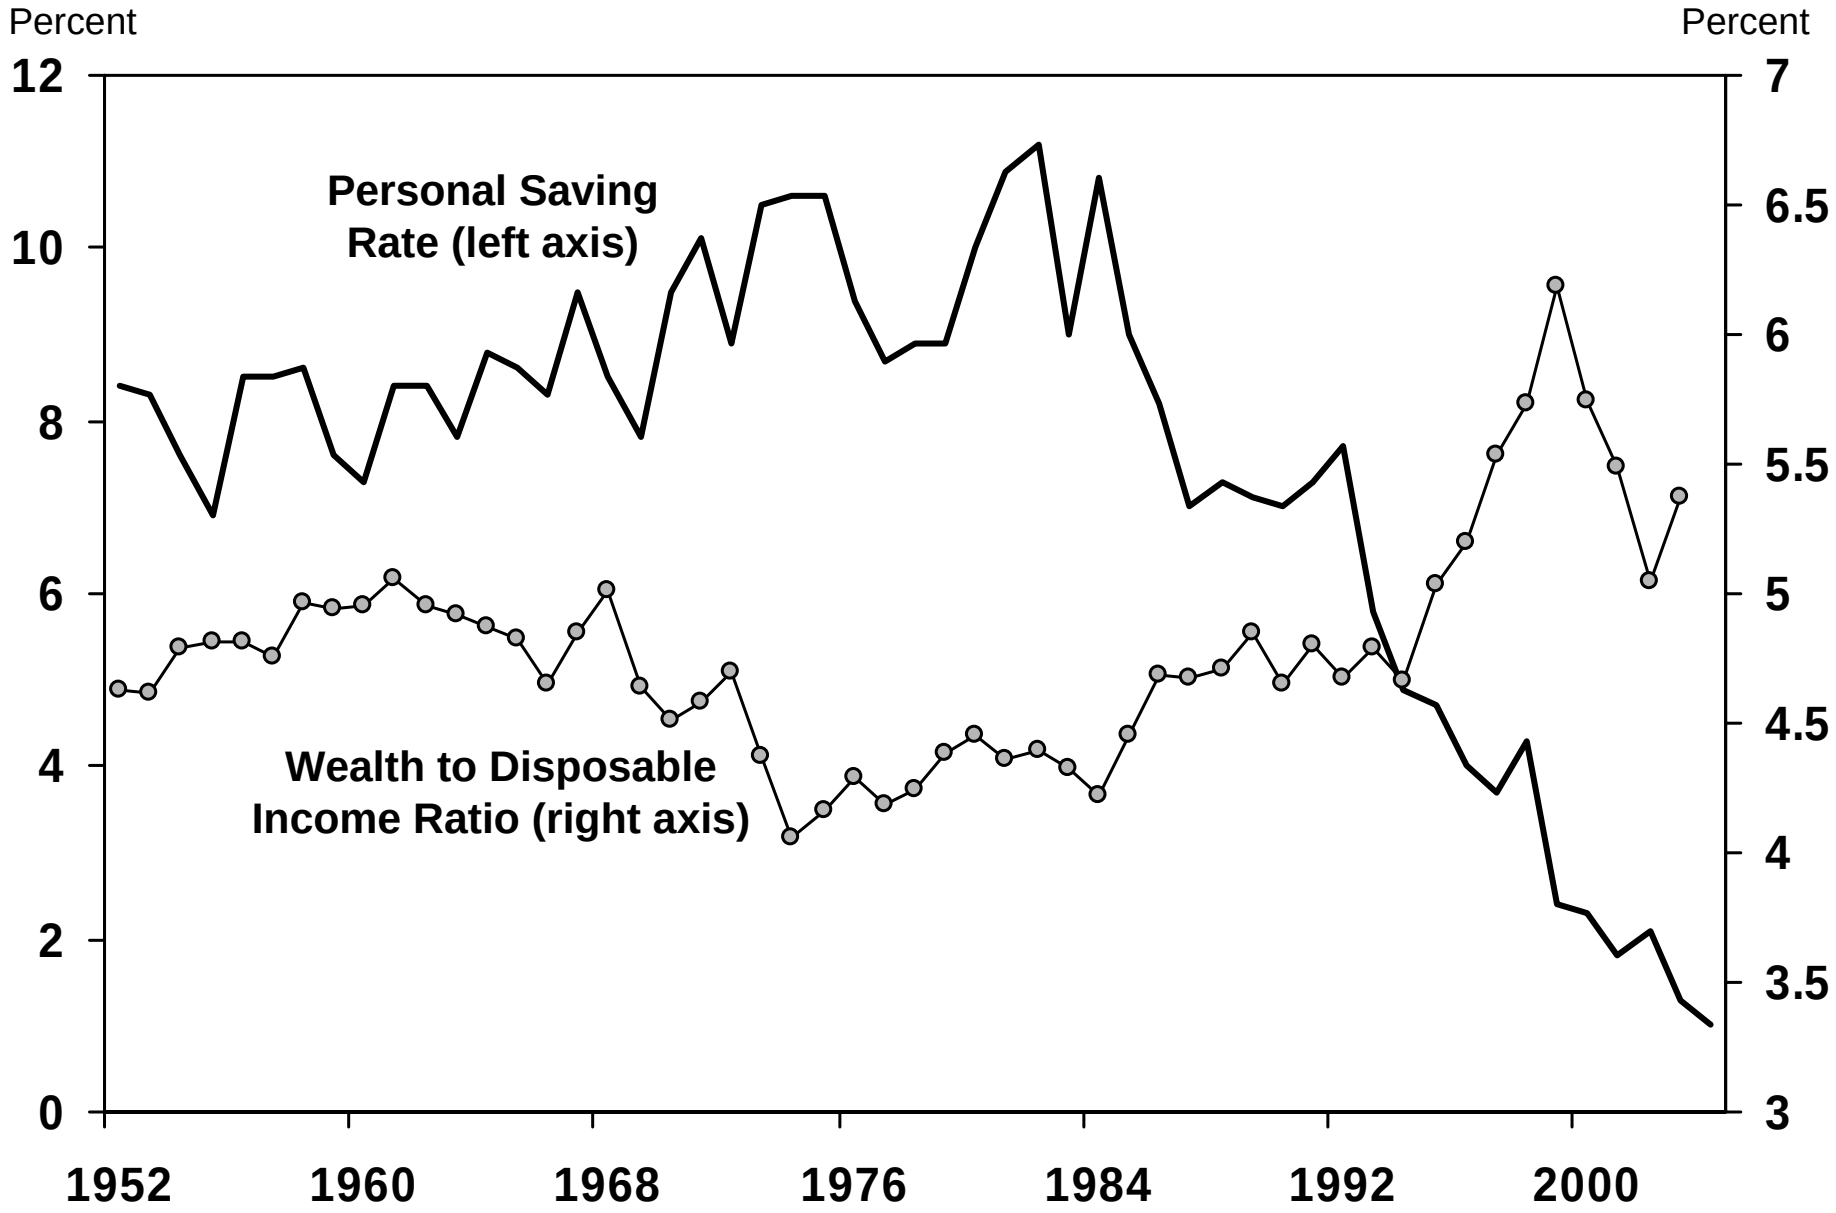

Source: Federal Reserve Board, Bureau of Economic Analysis

Chart 2: Real Wealth and Consumer Spending

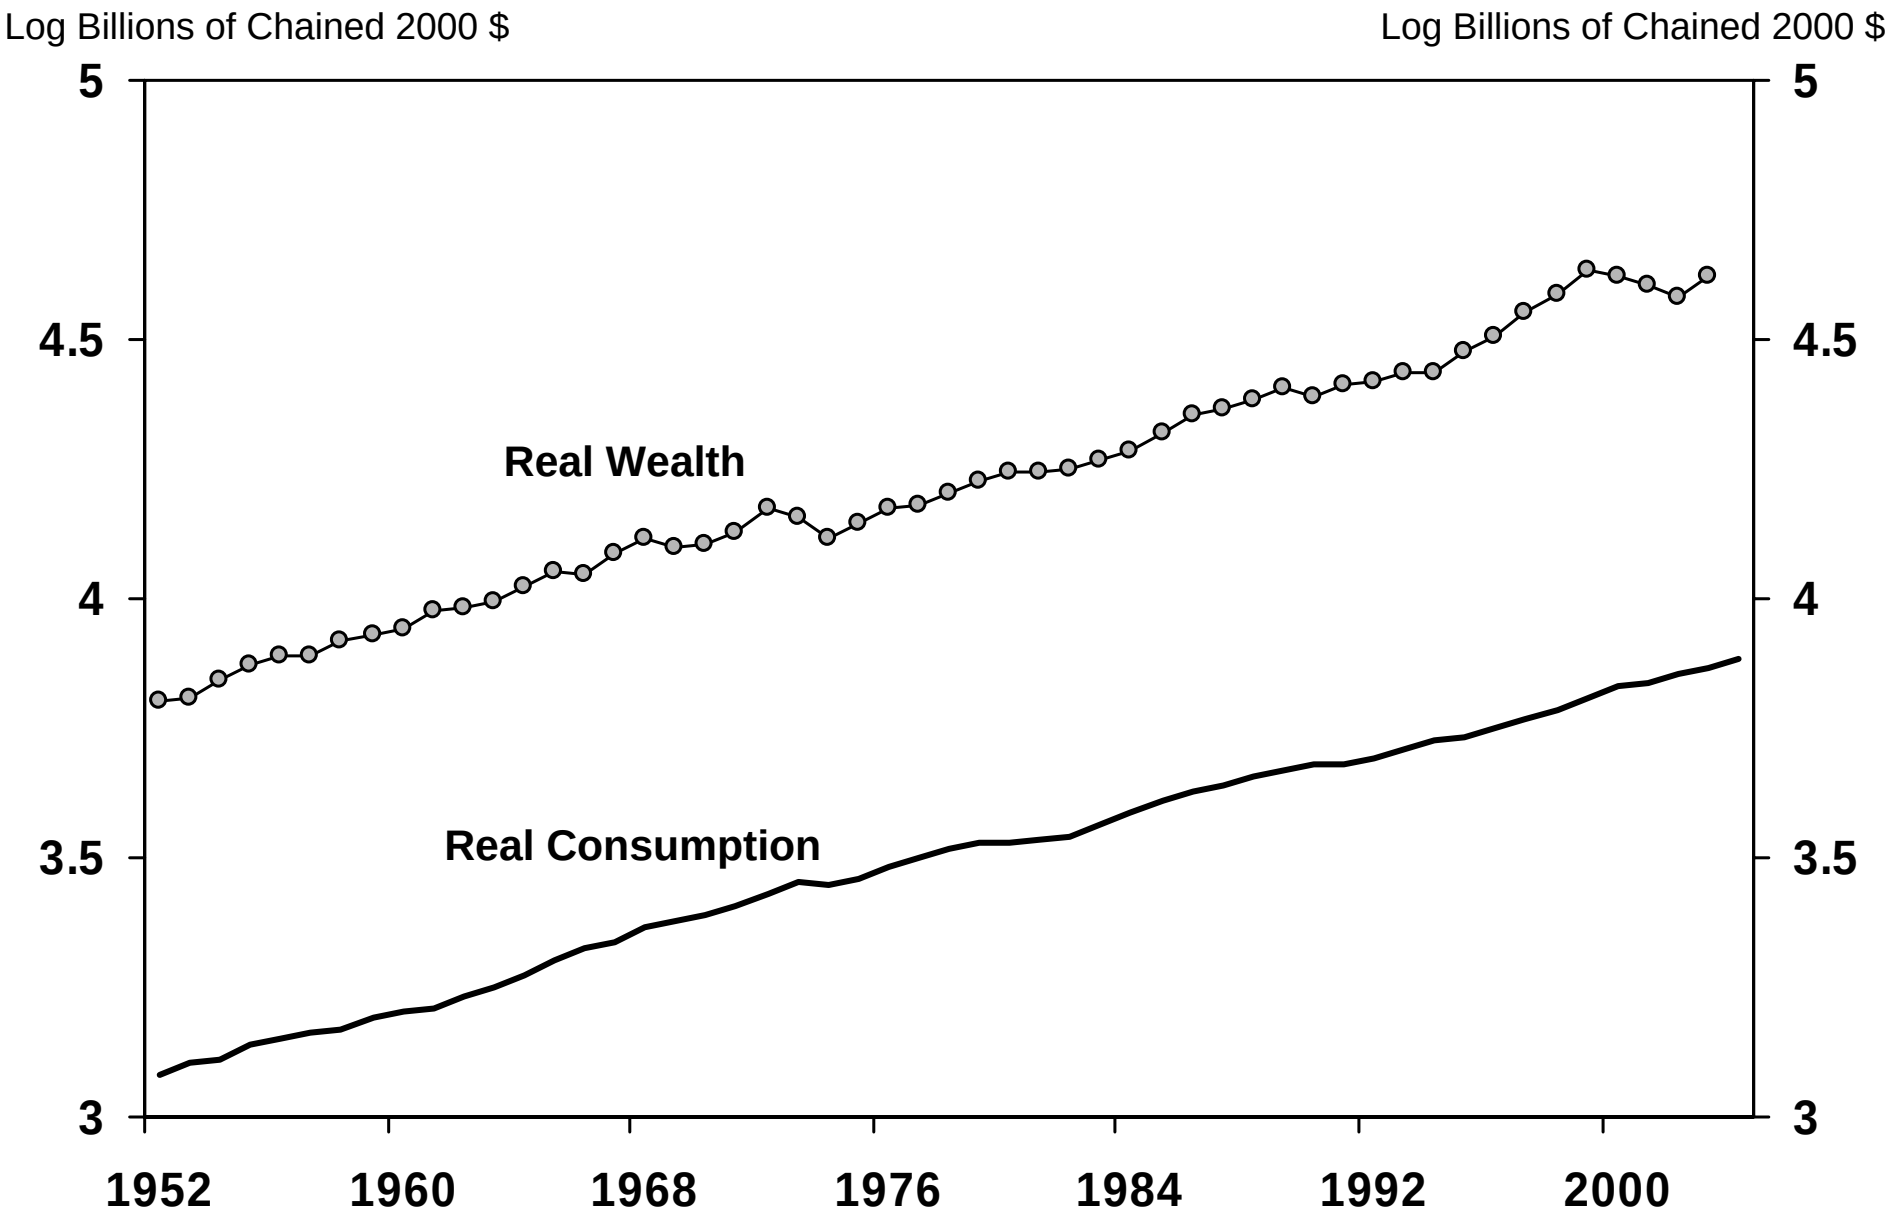

Source: Federal Reserve Board, Bureau of Economic Analysis

Chart 3: Decomposition of Changes in Wealth

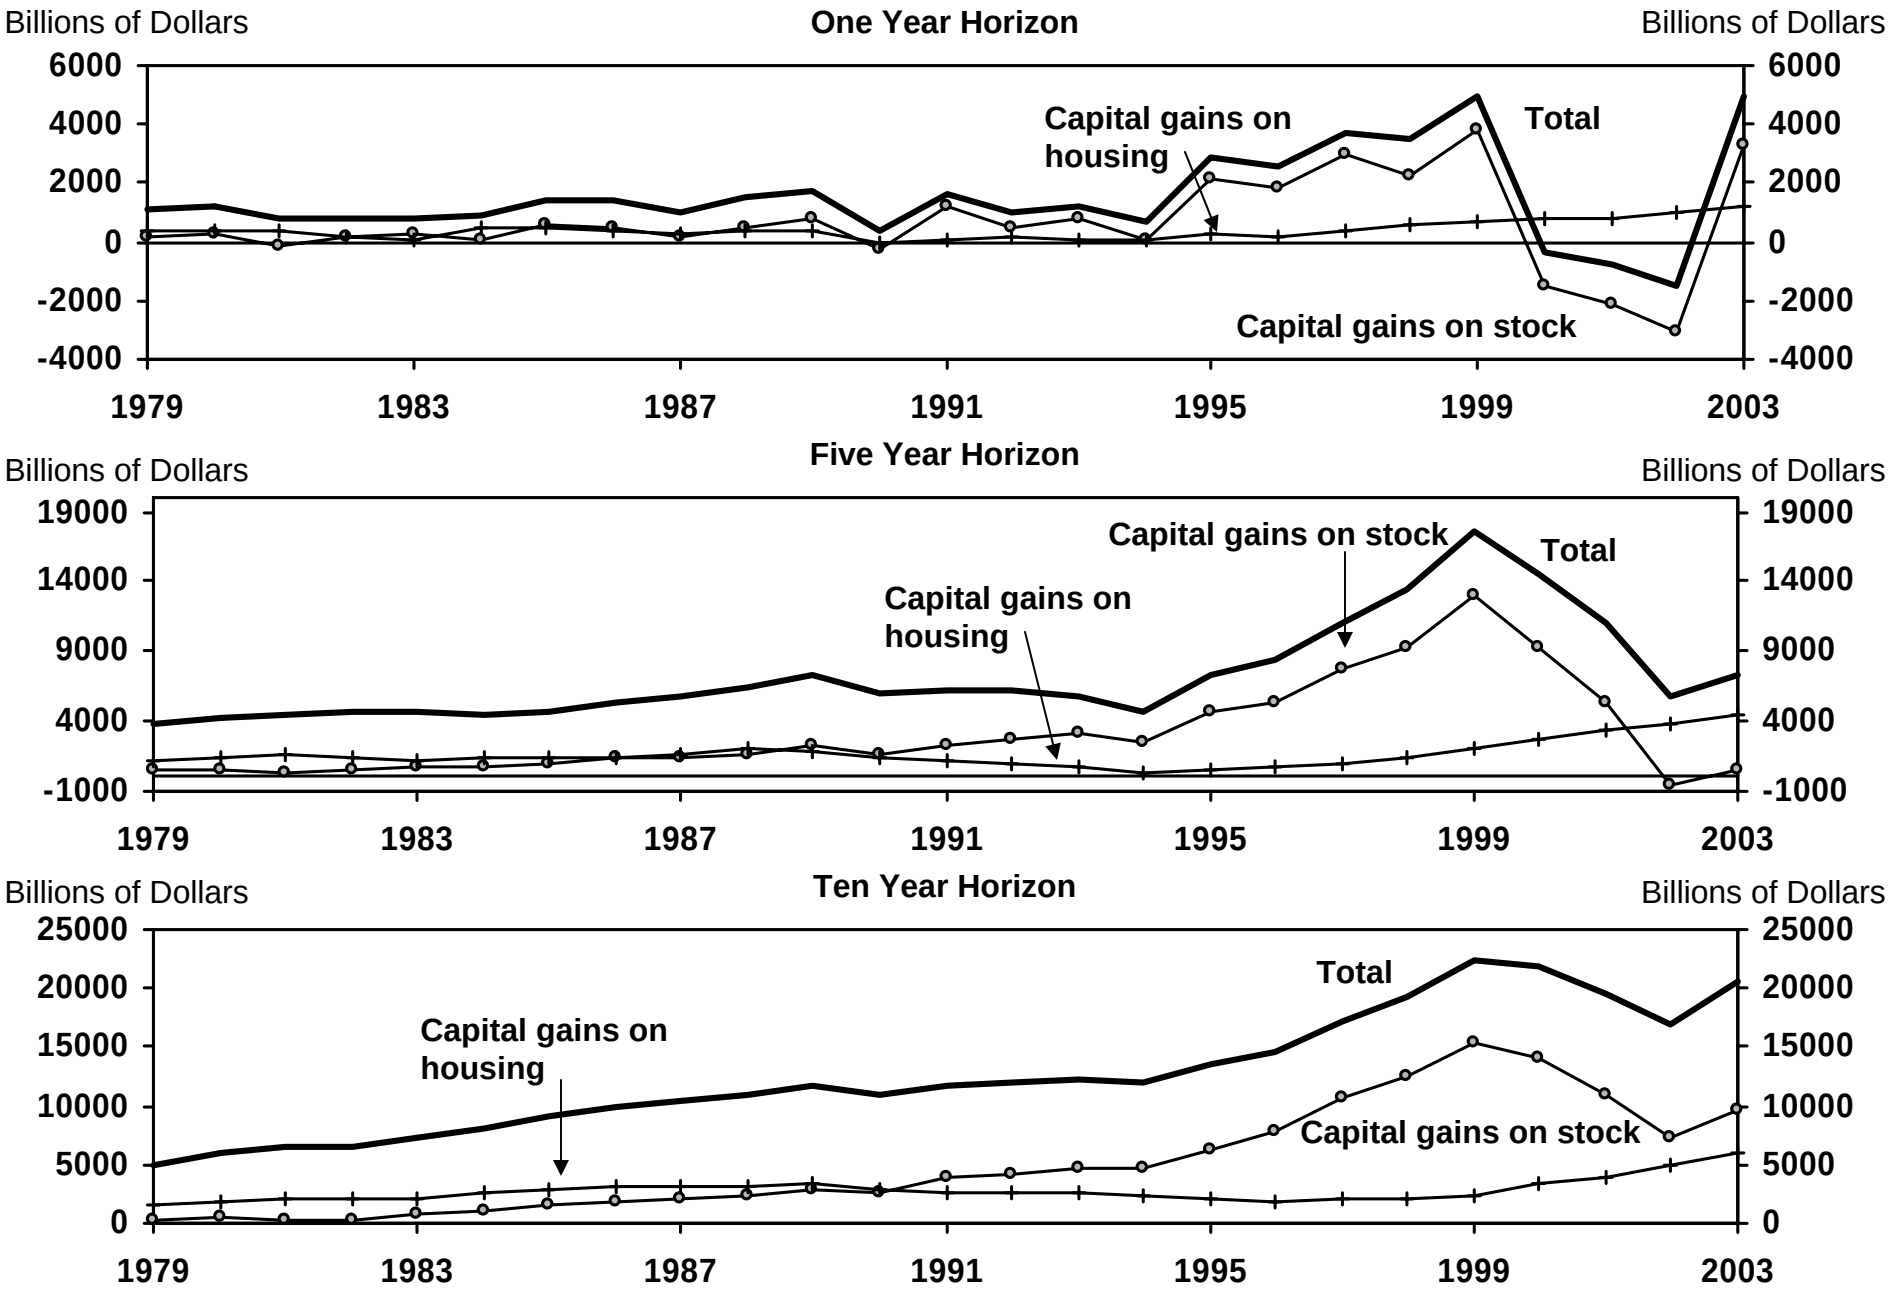

Source: Federal Reserve Board, author's calculations

Chart 4: Propensities to Consume From Wealth

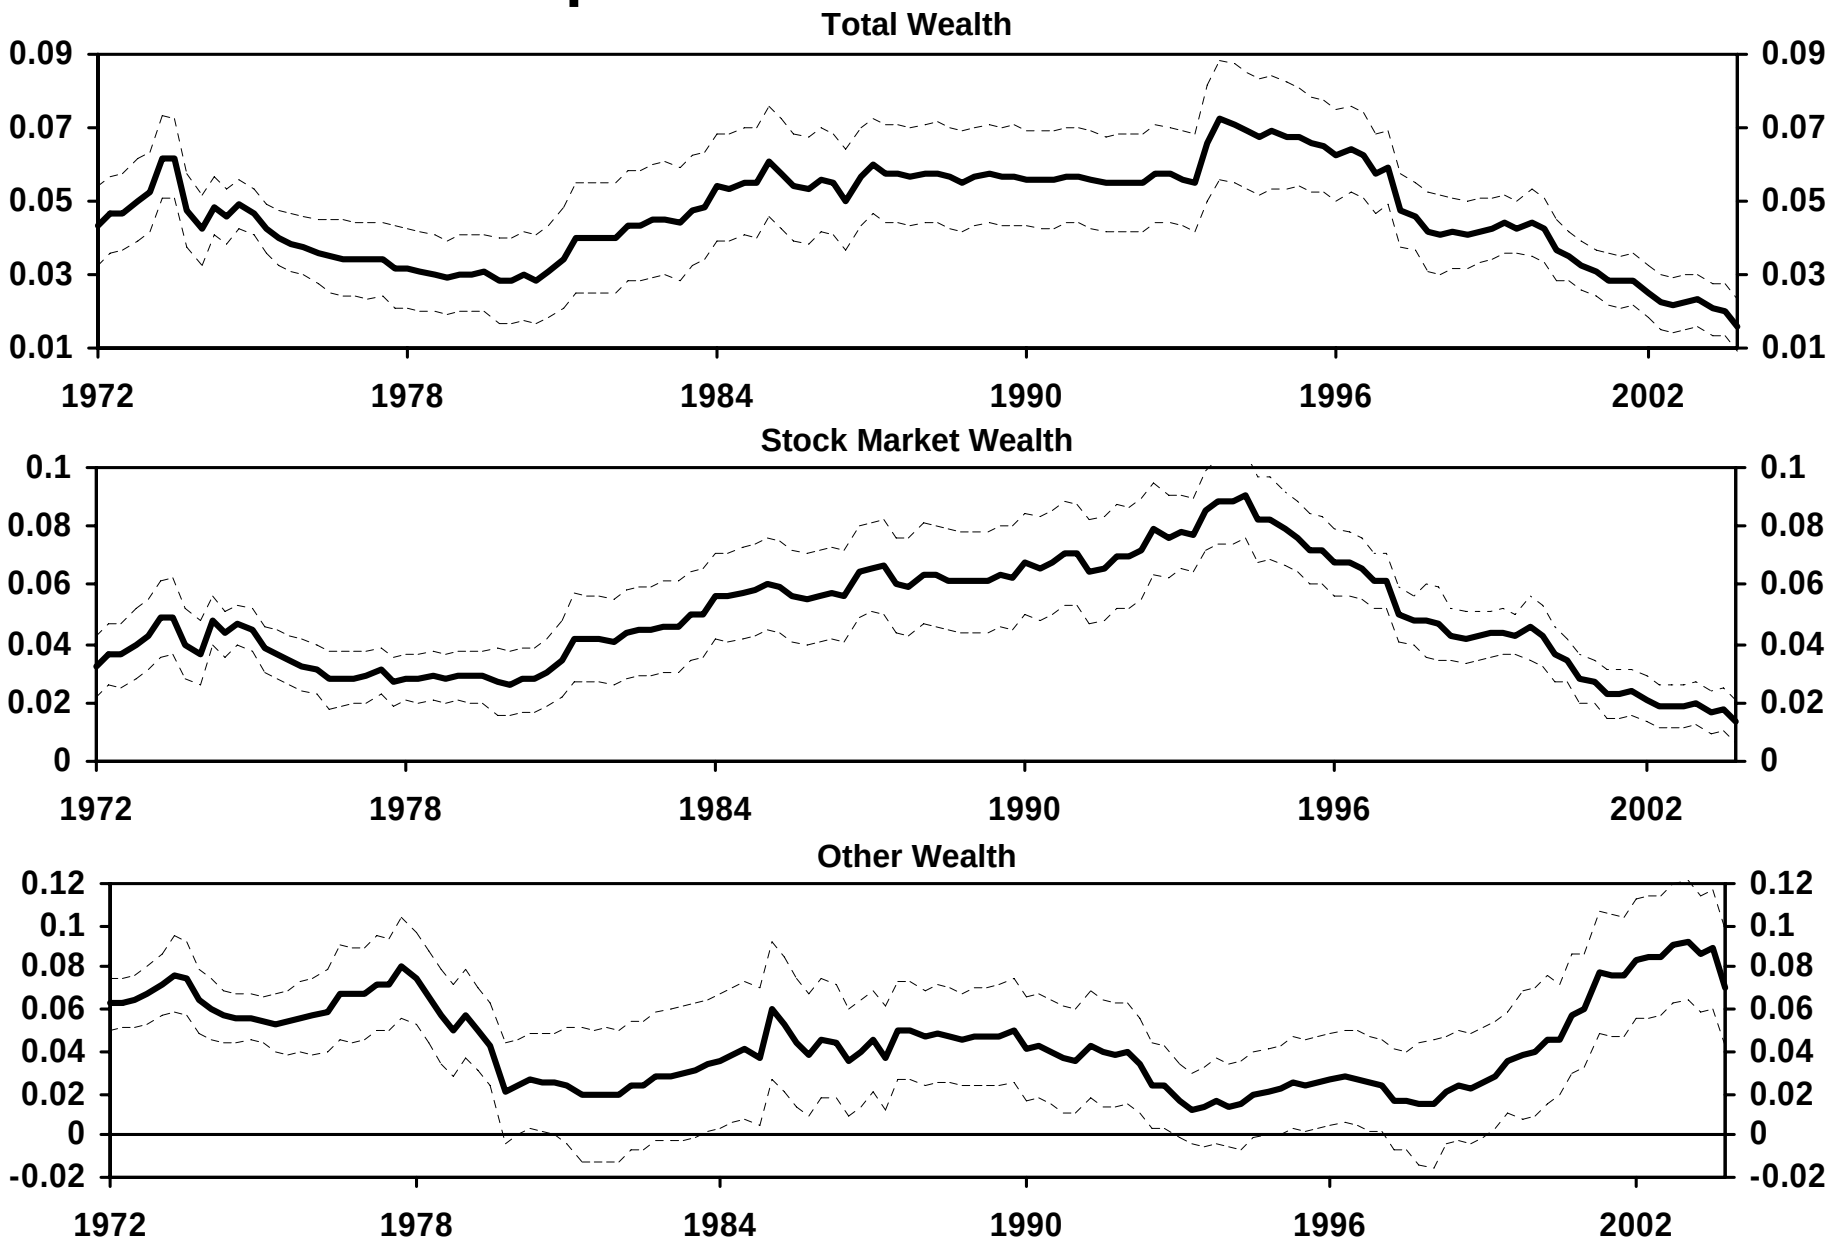

Source: Federal Reserve Board, Bureau of Economics Analysis, author's calculations

Results from 20-year rolling regression with enddate marked. Dashed lines represent one standard error around regression coefficients.

Chart 5: Per Capita Wealth

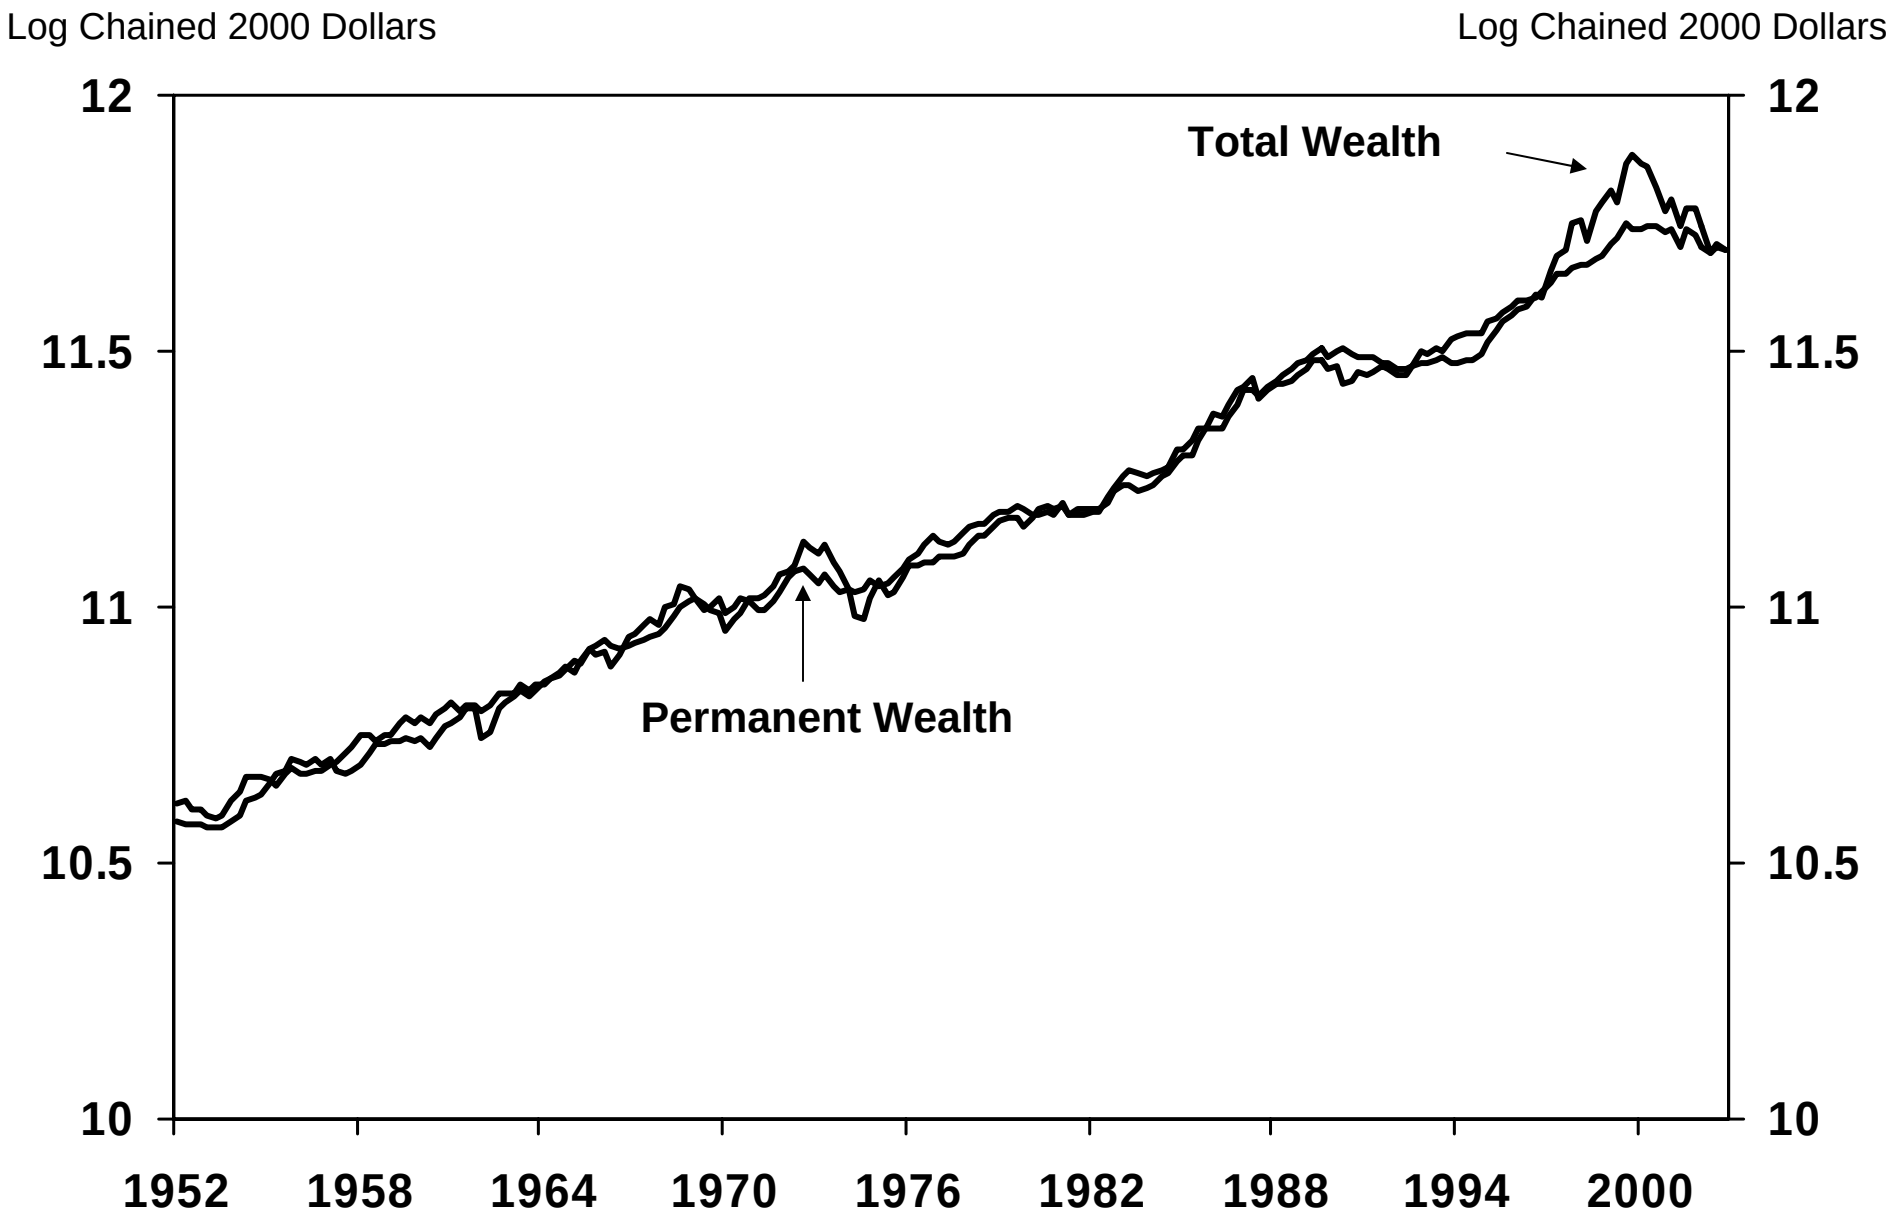

Source: Martin Lettau

Chart 6: Growth of Wealth, Income, Consumption

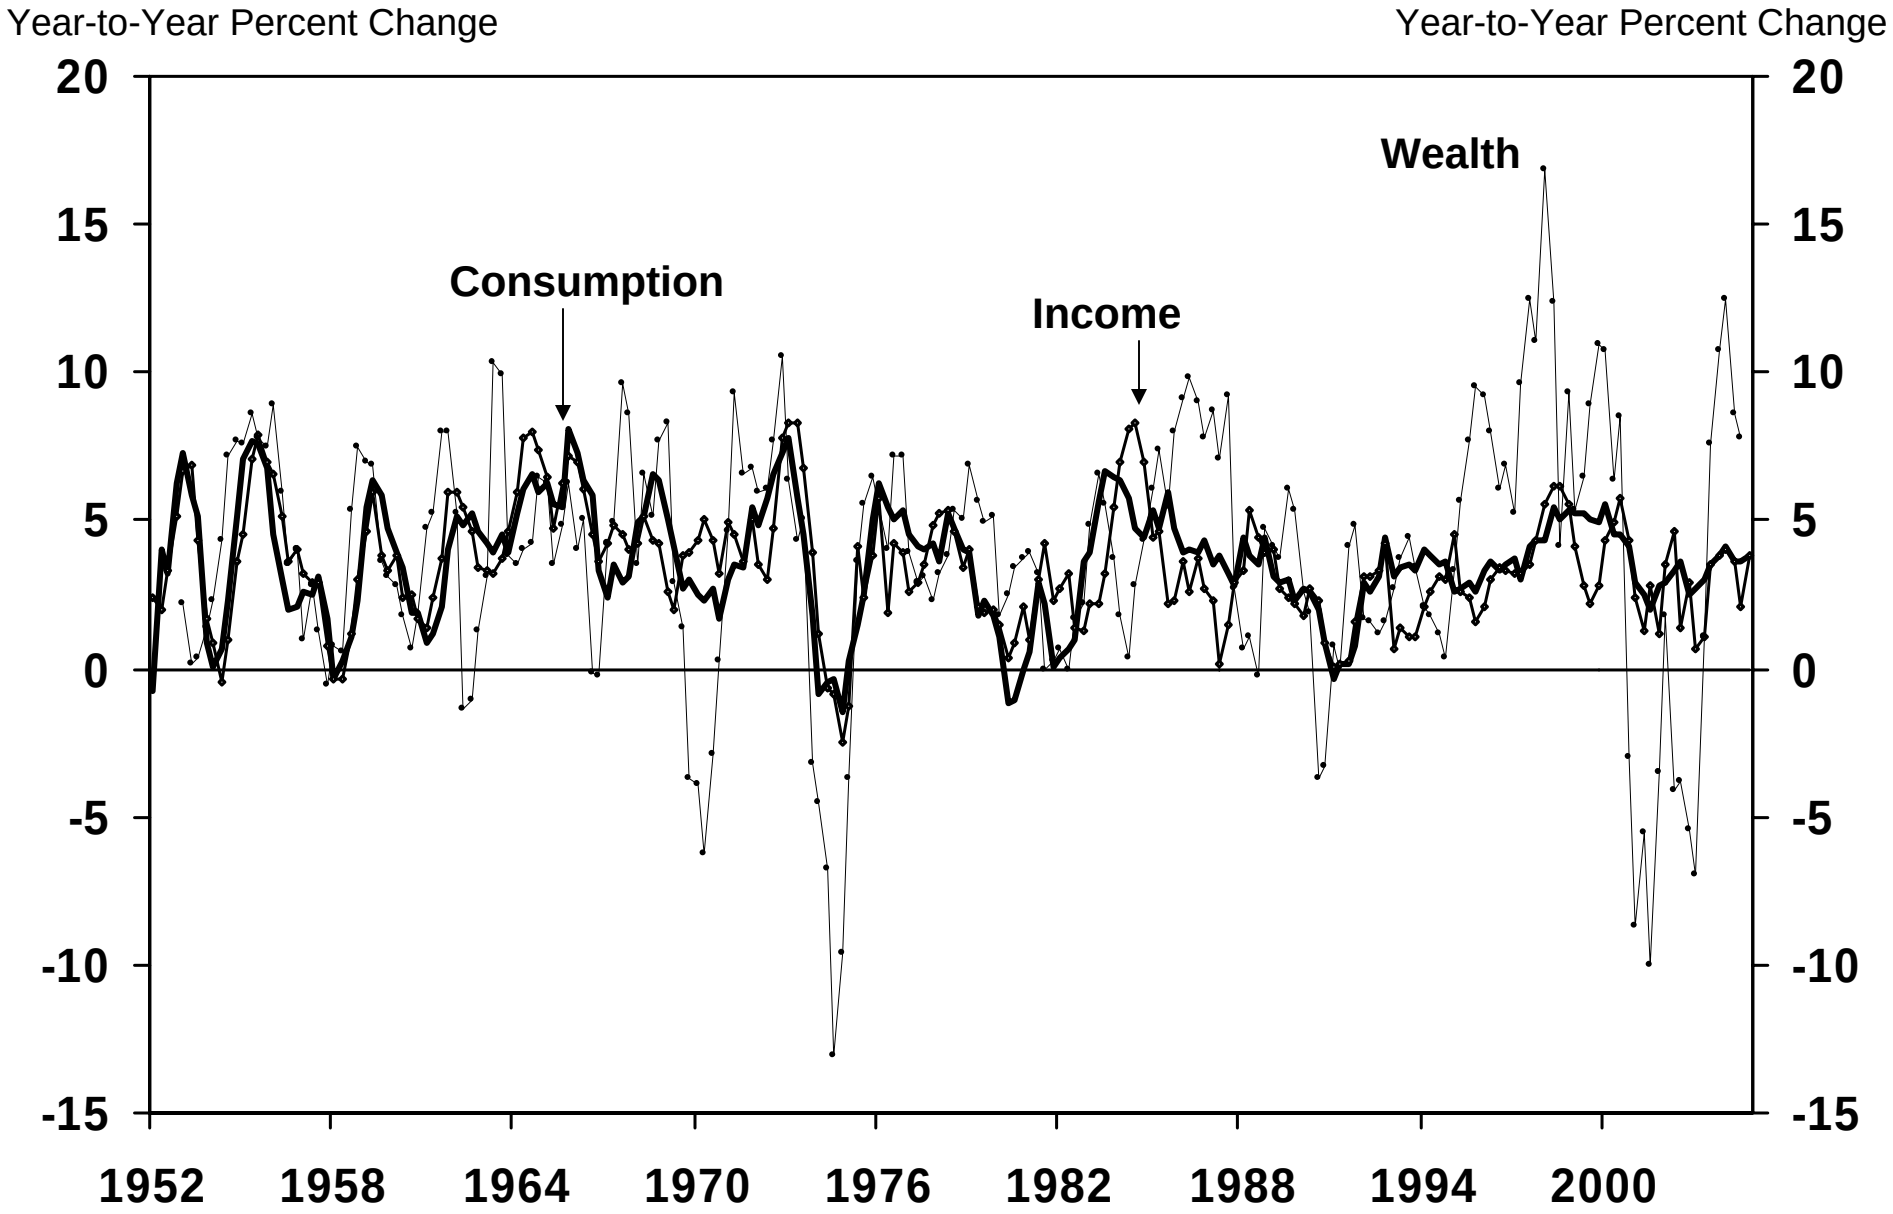

Source: Federal Reserve Board, Bureau of Economic Analysis