



WORKING PAPERS

RESEARCH DEPARTMENT

**WORKING PAPER NO. 07-8
MISMEASURED PERSONAL SAVING AND THE
PERMANENT INCOME HYPOTHESIS**

Leonard I. Nakamura and Tom Stark
Federal Reserve Bank of Philadelphia

First Draft: March 2005
This Draft: February 2007

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/econ/index.html

WORKING PAPER NO. 07-8
MISMEASURED PERSONAL SAVING AND THE
PERMANENT INCOME HYPOTHESIS

Leonard I. Nakamura and Tom Stark
Federal Reserve Bank of Philadelphia

First Draft: March 2005
This Draft: February 2007

JEL codes: E01 E21 C82

Keywords: Saving; Permanent Income; Real-time data

The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or of the Federal Reserve System. We thank Dean Croushore, Mike Dotsey, and Todd Clark for ideas and advice as well as participants at the Econometric Society 2006 Winter Meetings. We would like to thank John Chew, Mike Conlow, Will Olney, and Daniel Tannenbaum for superb research assistance. This paper is available free of charge at: www.philadelphiafed.org/econ/wps/index.html.

Mismeasured Personal Saving and the Permanent Income Hypothesis

Leonard I. Nakamura and Tom Stark

ABSTRACT

Is it possible to forecast using poorly measured data? According to the permanent income hypothesis, a low personal saving rate should predict rising future income (Campbell, 1987). However, the U.S. personal saving rate is initially poorly measured and has been repeatedly revised upward in benchmark revisions. We use both conventional and real-time estimates of the personal saving rate in vector autoregressions to forecast real disposable income; using the level of the personal saving rate in real time would have almost invariably made forecasts worse, but first differences of the personal saving rate are predictive. We also test the lay hypothesis that a low personal saving rate has implications for consumption growth and find no evidence of forecasting ability.

Correspondence to:

Leonard I. Nakamura, Research Department, Federal Reserve Bank of Philadelphia, 10 Independence Mall, Philadelphia, PA 19106, 215-574-3804 (office), 215-574-4303 (fax), leonard.nakamura@phil.frb.org (e-mail). Tom Stark, Research Department, Federal Reserve Bank of Philadelphia, 10 Independence Mall, Philadelphia, PA 19106, 215-574-6436 (office), 215-574-4303 (fax), tom.stark@phil.frb.org (e-mail).

I. Introduction

Recent years have witnessed increased interest in examining how revisions to the data affect empirical tests of economic hypotheses. One such test involves a theory's ability to predict. In this paper, we study an important economic hypothesis whose forecasting ability depends on a data series that is subject to large revisions. If consumption obeys the permanent income hypothesis, then as Campbell (1987) has shown, a low personal saving rate implies that real labor income is expected to accelerate. On the other hand, over the past four decades, whenever the personal saving rate is first published, it has almost always been too low and has been revised upward. We show that in real time, the level of the personal savings rate is uninformative for forecasting real disposable income growth—it is simply too noisy initially.¹ But in this case, real-time data permit us to go beyond this merely negative conclusion to a positive one. The first difference of the savings rate, we shall show, does have value in forecasting real income growth in real time.

We also investigate the persistent lay hypothesis that a low personal saving rate signals an overextended consumer and future decline in consumption, a hypothesis that has been put forward in U.S. monetary policy discussions.² We show that the level of the personal saving rate has no forecasting power for personal consumption expenditure, either in data that are revised or unrevised.

Economists have questioned how well personal saving is measured at least since Taubman (1968). Initially published estimates of the personal saving rate from 1965 Q3 to 1999

¹ We forecast real disposable income growth rather than real labor income growth because of data availability.

² For example, the minutes to the September 2004 meeting of the U.S. Federal Open Market Committee state, "Members perceived several possible sources of downside risk to household spending. In particular, households might hold back on spending in an attempt to increase their saving, which had fallen to a very low level relative to income."

Q2 have been revised upward more than 50 percent, from 5.3 percent to 8.1 percent, as we document. Most of these revisions are due to the benchmark revisions that follow economic censuses, with large revisions decades after the initial estimate, in turn due to large upward revisions to both disposable personal income and personal outlays. Nominal disposable personal income from 1965 Q3 to 1999 Q2 has been revised up 8.3 percent from initial publication.³ Benchmark revisions substantially change the relative ranking of saving rates for individual quarters and five-year averages. For example, the early 1980s are now viewed as the period with the highest saving rates in the postwar period; yet when the rates for that period were first published, they were reported to be the lowest saving rates since the Korean War.

Revision does appear to bring us closer to the true state of affairs that economic agents confront. For tests of the permanent income hypothesis carried out by Campbell (1987) and Ireland (1995), revised data are preferable. But revised data do not fit the informational situation of decision makers; out-of-sample forecasts using revised data are not good tests of likely real-time forecast performance (Croushore, 2006).

Real-time data have most often been used to check robustness, such as in Cole (1969), Diebold and Rudebusch (1991), and Stark and Croushore (2002), providing many additional sets of data with which to test economic theories. Some other work – such as Howrey (1978) and Koenig, et al., 2003 -- has focused on using real-time data to actually improve forecasts, as we do. These other efforts have focused on using the Kalman filter to improve forecasts. It is possible that this alternative approach could be fruitfully applied to our problem, but we pursue a simpler approach.

In what follows, we briefly review our data set and the process the U.S. Bureau of Economic Analysis (BEA) uses to revise national accounts data and show that the major changes to the personal saving rate have occurred in benchmark revisions. We then use both conventional

³ Boskin (2000) has pointed out the large upward revisions of nominal national income that have occurred over time.

and real-time estimates of the personal saving rate to forecast real disposable income. We confirm that adding the level of the personal saving rate to univariate AR models estimated on sufficiently revised data improves forecast power by conventional measures, just as the permanent income theory would suggest; we show in this case that the first difference of the saving rate provides inferior forecast power, as the theory would suggest. In real time, however, adding the level of the personal saving rate almost invariably makes the real-time forecasts worse, but adding the first difference of the saving rate improves them.

Our results thus highlight the important role of real-time data in designing forecasting models. First, we are able to avoid an empirical relationship based on the level of the saving rate that does not forecast well in real time. Second, we are able to discover an alternative that rescues some of the valuable forecast information the personal saving rate contains.

II. Real-Time Data and Revisions to the Personal Saving Rate

The Federal Reserve Bank of Philadelphia maintains a real-time data set for macroeconomists that consists of vintage snapshots of data as they were reported in the middle of each quarter from 1965 Q3 to the present; it is documented in Croushore and Stark (2001) and online at <http://www.philadelphiafed.org/econ/forecast/reaindex.html>.

Definitions and measurement difficulties. The personal saving rate is personal saving as a percentage of disposable (after-tax) personal income. Personal saving, in turn, is disposable personal income less personal outlays. Disposable personal income includes some easily measured items, such as social insurance benefits and contributions. Other parts of labor income, such as other benefits and transfers, are subject to measurement and conceptual problems. Wages and proprietors' income are subject to underreporting in government records as a result of tax evasion. And rental income and proprietors' income are net income measures that require estimates of depreciation and other expenses that are hard to measure well. Capital gains on equity (other than from qualified equity stock options) and real estate are not included in personal income.

A general rule of national income accounting is to ignore income from capital gains, whether realized or not. BEA has chosen to include realized capital gains from employee stock options in its measures of personal income. (These capital gains are subtracted from corporate profit, so there is no net impact on gross domestic income.) Real capital gains, measured by the increase in stock market value of domestic corporations, averaged 10.4 percent of real disposable income from 1984 to 2004 and only 1.9 percent from 1954 to 1984. Thus personal income may be understated to the extent that the returns from equity holdings appear as (uncounted) capital gains rather than (counted) dividends and employee stock options and to the extent that the rental return to property ownership omits the capital gains from rising house prices.⁴ Personal outlays -- personal consumption expenditures (95 percent of personal outlays) plus transfers and nonmortgage interest payments⁵ -- are generally easier to measure as we explain below.

The data revision process. The BEA revises the national income accounts as follows. Data on a given quarter's economic activity are first published in an *advance* estimate, late in the first month of the next quarter.⁶ The data available at this time are recorded in the Philadelphia Fed's real-time data set as the vintage of that quarter. The *revised* estimate is published in the second month of a quarter followed a month later by a *final* estimate. These data are then generally left unchanged until the following summer, when the latest three years of national account data are revised.⁷ A set of initial estimates thus undergoes three *summer* revisions. Thereafter, the estimates are changed only in what are called *benchmark* revisions, which now occur every four years. Benchmark revisions provide the opportunity for BEA to make

⁴ Rental income, including implied income from owner-occupied housing, in 2004 was \$166 billion. This is a 1.2 percent nominal return on net equity of housing (for households, nonprofits, and nonfarm, noncorporate businesses) of \$13.7 trillion from the Federal Reserve's flow of funds data. Over the entire period from 1965 to 2004, according to latest vintage information, the average nominal return was 1.5 percent. At the same time, the return to the 12-month constant-maturity U.S. Treasury bill averaged 6.6 percent.

⁵ Mortgage interest payments are netted out of rental income.

⁶ Until 1985, the BEA also published a "flash" GDP estimate 15 days before the end of a quarter, but this estimate included only aggregate nominal and real GDP, without any underlying detail (although some detail was circulated internally within the government), and did not include the personal saving rate.

⁷ One change in the routine has been that wages and salaries, since 2002 Q3, are revised again three months after the final estimate.

discretionary choices in defining the items it considers to be part of personal income; in addition more complete data from economic censuses are included at this time.

The personal saving rate from 1965 Q3 to 1999 Q2 was 5.3 percent if averaged over the rate first observed in the advance estimates (Figure 1); by the 2005 Q3 vintage it averaged 8.1 percent. Thus the personal saving rate over time has been revised systematically upward.

The upward revisions occur in benchmark revisions. As Figure 2 shows, revisions that occur between the advance estimate and the last vintage before any benchmark revision have been relatively unbiased and small, with the mean rise of 0.08 percentage point and a mean squared revision of 1.11 percentage points.⁸ By contrast, the revisions from advance estimates to the latest vintage (the data published in 2005 Q3) have a mean of 2.44 percentage points and a mean squared revision of 9.52 percentage points (Figure 3). The benchmark revisions thus account for very nearly all of the bias and the bulk of the mean squared revision.

Upward revisions to disposable personal income have been very large. Disposable personal income has been revised as much as 14.8 percent; on average, from 1965 Q3 to 1999 Q2, the revision has been 8.4 percent (Figure 4). Over the same period, nominal GDP and personal outlays were revised up by less: nominal GDP by 6.5 percent and personal outlays by 5.1 percent.

These large increases are the result of steady upward shifts. Our data begin with the observation for 1965 Q3, as recorded in the vintage of 1965 Q4, and they are averaged into five-year periods (Table 1). Of the 26 changes that these groups underwent in benchmark revisions, 16 were positive and greater than 0.5 percentage point. Another view of these large increases can be seen in the vintage data presented in Figures 5a – 5c. In each of these nine vintages, spread over the period from 1980 to 2005, the most recent saving rates were well below the average.

⁸ This figure and accompanying data omit the advance estimates that occur just before a benchmark revision, and thus had no opportunity to change. We also excluded the last advance estimate, for 2005 Q2, for the same reason from both this and the next figure.

The revisions in the first three years after the data are first published are primarily from regular sources whose availability is delayed. Systematic biases related to these data can be estimated and eliminated, and BEA apparently has done so. Benchmark revisions, on the other hand, incorporate two basic types of changes: statistical changes, based on newly available data, and definitional changes. Statistical changes include data from censuses, such as the economic census or the population census, and other sources of data that become available with a long lag or irregularly, such as IRS random audit data. Definitional changes include changes in data recognition (such as reclassification of government pension contributions as personal income) and changes in concept (such as including software as investment or introducing chain-weighted prices).

Why are saving rate revisions biased upward? Two factors drive revisions: income is harder to measure than expenditures and economic evolution creates new sources of income. Income is harder to measure than final expenditure because it must be collected from more units. The vast bulk of gross domestic product measured from the expenditure side is final sales of domestic purchasers. Final sales data – purchases by consumers, businesses, and governments – do not require information about the entire production chain, only the final point of sale. By contrast, to obtain gross domestic income we need data from each industry on labor and capital income. Retailers, for example, account for nearly one-third of final product but only about one-fifteenth of labor income. Since data on income are costlier to collect, more of it escapes counting initially. Income-side data are aggregated to gross domestic income (GDI), conceptually the same as GDP, but in practice differing by between 2.3 percent and -1.8 percent; GDP minus GDI is called the statistical discrepancy. Generally speaking, the statistical discrepancy is positive – since 1965 Q3 it has averaged 0.7 percent, with 131 of 160 observations positive – suggesting that typically income is undercounted. The fact that GDP and not GDI is used as the primary yardstick expresses BEA's judgment that it is the more precisely measured of the two aggregates.

As the economy evolves, new types of income and expenditures – such as stock equity options, software, and Internet sales – arise. Initially these may not be included in BEA’s national income measures. Over time, comprehensive measures of economic activity, such as economic censuses and tax audits and reassessments of income definitions, incorporated in the benchmark revisions, will tend to expand the universe to new industries and practices. To the extent that more income has been missed because it is harder to measure, saving will rise over the course of successive revisions.

John Campbell (1987) has argued that labor income can be specified as a unit root process. Saving, in the theory he presents, represents expectations of future declines in income and is stationary in its level. However, we have argued that saving is not measured accurately. Some fraction of income is likely undercounted in the initial estimates. It is reasonable to think that once the missing income is discovered by national income accountants and incorporated into their subsequent estimates, income, and hence saving, will be subject to permanent revisions in the same direction. These revisions could have a large effect on the real-time forecasting performance of any model that hinges directly on the saving rate as its driving variable.

Indeed, the missing income that is recognized in benchmark revisions typically has a unit root. Define a benchmark revision as $r_{st}^{(B)} = s_t^{(B)} - s_t^{(B-1)}$, $t = 1 \dots B - 2$. Here s_t is the personal saving rate at date t , as recorded at date $B-1$, where B is the date of a benchmark revision. For benchmark revisions, the hypothesis of a unit root can generally not be rejected. Considering the benchmark revisions of 1976, 1981, 1986, 1991, 1996, and 1999, only for the 1996 revision can the hypothesis of a unit root be rejected at the 10 percent level. It appears that the missing income and saving recognized in benchmark revisions typically has a unit root.

If $r_{st}^{(B)}$ has a unit root, then even though the revised saving series, $s_t^{(B)}$ does not have a unit root, the pre-revision saving series, $s_t^{(B-1)}$, will have a unit root (although in practice the unit

root may not be empirically detectable). This in turn may have serious consequences for the econometric properties of saving data before they are revised.

III. How Closely Related Are the Advance and Revised Estimates of the Saving Rate?

If initial and revised estimates of the saving rate were highly correlated, revisions would likely have little effect on empirical tests of economic hypotheses. The collection of real-time data sets would be relatively unimportant, and we could be comfortable in assuming that any results obtained from the latest vintage of data would also hold in real time. From a forecasting perspective, we would expect performance based on the latest vintage data to roughly match the performance we would achieve in real time. Indeed, such an assumption underlies the important work of Campbell (1987) and Ireland (1995).

To begin our analysis of the relationship between initial and revised estimates, consider a regression whose left-hand-side variable is a given vintage (V) personal saving rate (PSR_V) and whose right-hand-side variables are the original advance estimate of the personal saving rate (PSR_A) and a constant:

$$PSR_{V,t} = \alpha + \beta PSR_{A,t} + e_t. \quad (1)$$

Mankiw and Shapiro (1986) use such an equation to test whether the advance estimate is an efficient forecast of the revised estimate. Under the joint null hypothesis that $\alpha = 0$ and $\beta = 1$, the advance estimate is said to be an optimal forecast of the revised value. We show the results of this test below. From a forecasting perspective, however, we are more interested in the correlation between the advance and revised estimates. We thus concentrate on the regression equation's R-squared.

The results of estimating equation 1 for overlapping 20-year periods are shown in Table 2A (we choose 20 years because we can get two complete nonoverlapping groups into our 40-

year sample), taking five-year intervals for our analysis. The most telling results are those in the first row. For the period from 1965 Q3 to 1985 Q2, the estimated slope is, for all periods, significantly different from unity, using Newey-West HAC robust standard errors. More important, the unadjusted R-squared declines with successive vintages. In the R-squared results for the 2005 Q3 vintage, the advance estimates account for at most 43 percent of the variance of the revised estimates and as little as 12 percent.⁹

How are the first differences of the personal saving rate (DPSR) influenced by the revision process? Table 2B shows the result of substituting first differences of the personal saving rate in place of the level in the regressions. Note that the coefficient on the first difference of the advance estimate does not fall appreciably over successive benchmarks.¹⁰ And the early vintage variations in the first differences of the personal saving rate capture a large part of the variation in later vintages. The R-squared statistics are remarkably stable across different vintages of the left-hand-side variable.

Our conclusion from this regression analysis is that early estimates of the level of the saving rate are not closely tied to the revised values. Indeed, the correlation drops as the data are changed in one benchmark revision to the next. It is possible that a large part of the current difference between the high saving rates in the early 1980s and the current low saving rate will also prove to be the result of measurement error. What is clear is that these data are measured with considerable noise, and there is little reason to believe that our measures have become more stable than in the past. However, the same is not true of first differences of the personal saving rate, which tend to be affected much less by benchmark revisions. In the following section, we quantify the effect of revisions on the forecasting performance of models that rely on the saving rate as a key predictor.

⁹ Throughout, the personal saving rate fails the Mankiw et al. test.

¹⁰ For the advance estimate, the first difference is taken within each vintage. It is, in other words, the advance estimate of the first difference. For example, for the first differenced 1984 Q1 advance estimate, the 1983 Q4 saving rate is subtracted from the 1984 Q1 saving rate, both taken from the 1984 Q2 vintage.

IV. Using the Personal Saving Rate to Forecast in Real Time

We now address the question of whether the personal saving rate is too noisy in practice to be useful in forecasting in real time. If saving rates are low, should we expect that future income will rise relative to consumption as saving rates mean-revert? Or, as the lay hypothesis suggests, will consumption fall?

The lay hypothesis of the overextended consumer. First, let us address the lay hypothesis that a low personal saving rate implies a future decline in the growth rate of real personal consumption expenditure. A persistent lay belief is that a low personal saving rate is indicative of households being overextended and portends lower real personal consumption expenditures in the future. Although the pure permanent income hypothesis implies that real consumption growth can not be forecast, that does not preclude the personal saving rate's having forecast ability for personal consumption expenditures, if, for example, a low personal saving rate implies fewer purchases of consumer durables in future periods, due to credit constraints.

We compare the root-mean-square error (RMSE) of a univariate autoregressive forecast, in which lags of the quarter-over-quarter percent change in real personal consumption (Δc_t) are used to forecast real personal consumption growth, to the RMSE of a bivariate vector autogression (VAR), which adds lags of the level of the personal saving rate (s_t). Our estimations use an expanding window of observations, adding an additional observation as we roll through each quarter of our sample. To compute forecast errors, we use three methods. First, we use latest vintage data (the data available in 2005 Q3) to estimate and forecast the models and to construct the forecast errors. We call these forecasts LV. This addresses the theoretical relationship as revealed in revised data. We then use real-time data to estimate and forecast, computing forecast errors in two ways. We first compute forecast errors based on real-time realizations, to see whether the data help predict consumption growth as reported at the time

(RT). This test shows the ability of forecasts to track data as revealed in the short run, as in the forecast contests that business economists are often judged by. We also compute forecast errors based on latest available vintage data (RTL_V). Since the latest vintage data have a tighter relationship to economic fundamentals (such as sales, unemployment, inflation, or interest rates), a decision maker might prefer this latter test. To specify lag length, we use, alternatively, the Akaike information criterion (AIC), the Schwarz information criterion (SIC), and a fixed number of lags, set to six. In the case of the AIC and SIC, lag length is re-estimated each quarter. We consider four forecast horizons: One-step-ahead quarter-over-quarter growth, two-step-ahead two-quarter average growth, four-step-ahead four-quarter average growth, and eight-step-ahead eight-quarter average growth. We analyze forecast performance over the period 1971 Q1 to 2005 Q2. For reasons that we discuss below, we also examine the subperiods 1971 Q1 to 1981 Q4 and 1982 Q1 to 2005 Q2.

Table 3 records the ratio between the root-mean-square errors of forecasts with the personal saving rate and without it; a ratio of less than one implies forecast improvement when the model includes the saving rate. We see that the level of the personal saving rate has no forecasting ability over the entire sample period (1971 Q1 to 2005 Q2), either with the latest vintage data or in real time. Nor does it have forecasting ability over either subperiod. A low saving rate does not predict future declines in consumption at any forecast horizon.

Permanent income hypothesis forecasts. Campbell (1987) has argued that low saving should be a signal of expected future growth in labor income. In a bivariate vector autoregression of saving and real labor income growth, lags of saving should have a negative sum, according to this theory, so that increases in saving forecast declines in real labor income. Campbell's regressions, covering 1953 to 1985, confirmed that high saving did forecast slower real labor income growth.

Campbell also showed that the permanent income hypothesis implied a tight set of cross-equation restrictions between the coefficients of the bivariate VAR. The intuition behind these

cross-equation restrictions is that a future predictable permanent increase in real labor income should generate a current permanent increase in consumption and therefore a temporary decrease in saving. When the permanent increase in real labor income arrives, the saving rate rises at the same time. Formally, the system is

$$\begin{bmatrix} \Delta Y_{it} \\ S_t \end{bmatrix} = \begin{bmatrix} a(L) & b(L) \\ c(L) & d(L) \end{bmatrix} \begin{bmatrix} \Delta Y_{it-1} \\ S_{t-1} \end{bmatrix} + \begin{bmatrix} u_{1t} \\ u_{2t} \end{bmatrix} \quad (2)$$

where Y_{it} is real labor income per capita at time t , S_t is real saving per capita, and Δ is the first-difference operator. The terms $a(L)$, $b(L)$, $c(L)$, and $d(L)$ are polynomials in the lag operator, given by, for example, $a(L) = \sum_{i=1}^p a_i L^i$, p is the lag length, and the u_i are forecast error terms.

The 2p restrictions on the coefficients of the lag operators are

$$\begin{aligned} c_i &= a_i, i = 1, \dots, p \\ d_1 &= b_1 + (1+r) \\ d_i &= b_i, i \geq 2, \end{aligned}$$

where r represents a constant real interest rate. Campbell finds these tight restrictions of the model are strongly rejected when more than one lag is included in the bivariate VAR.

Campbell's empirical work was revisited by Peter Ireland (1995). The coefficients on saving had a negative sign as Campbell's hypothesis predicts. Once again, the cross-equation restrictions were strongly rejected. But Ireland pointed out that, as noted by King (1995), formal hypotheses seldom fail to reject the implications of detailed mathematical models. A better test, Ireland argued, might be out-of-sample forecast performance. Using latest vintage data, Ireland tested the constrained VAR's and unconstrained VAR's rolling out-of-sample forecasts of one, two, four, and eight quarters ahead against the univariate forecast for real labor income. At forecast horizons of one, two and four quarters ahead, the unconstrained VAR improved on the univariate model of real labor income. In addition, the constrained VAR improved on the

univariate model at all forecast horizons and improved on the unconstrained VAR at all horizons except the one quarter ahead, where they tie. Ireland took this to be evidence in favor of the permanent income hypothesis.

In light of the behavior of the revisions to the personal saving rate, we wish to revisit these findings to see whether the forecasts would have been improved *in real time*. To do this, we estimate our model and make our out-of-sample forecasts using real-time data. We compute forecast errors based on real-time data (RT) and on latest available vintage data (LV).

Our data are not the same as Ireland's: we have real-time data on real disposable income but not on real labor income per capita, and we use the personal saving *rate*, not real saving per capita.¹¹ We also model the percent rate of growth in real disposable income, not the change. Labor income, a constructed variable that excludes dividend income, interest income, and the capital share of proprietor's income, is not a variable published as such by BEA. Our version of the VAR is thus given by

$$\begin{bmatrix} \Delta y_t \\ s_t \end{bmatrix} = \begin{bmatrix} a(L) & b(L) \\ c(L) & d(L) \end{bmatrix} \begin{bmatrix} \Delta y_{t-1} \\ s_{t-1} \end{bmatrix} + \begin{bmatrix} u_{1t} \\ u_{2t} \end{bmatrix},$$

where Δy_t is the first difference of the log of real disposable income in period t , and s_t is the personal saving rate, defined as nominal personal saving divided by nominal disposable income.¹² To see whether these substitutions create a large difference, we replicate Ireland's unconstrained analysis using the same sample period, vintage, and lag length, substituting the rate of growth in disposable personal income for the change in real labor income per capita and the personal saving rate for real personal saving per capita, in Table 4. This is to show that the

¹¹ All data used in this study are available on the Philadelphia Fed's web page. We deflate real-time observations on nominal disposable income with real-time observations on the personal consumption expenditure deflator.

¹² Each equation of the VAR also includes a constant.

essential features of the estimation are not disturbed by the inclusion of some of the capital income measures that Campbell and Ireland have excluded and by our use of the saving rate. As Ireland did, we use six lags and test the forecasts for one-step-ahead quarter-over-quarter growth, two-step-ahead two-quarter average growth, four-step-ahead four-quarter average growth, and eight-step-ahead eight-quarter average growth. If we focus on the latest vintage of data (LV) that Ireland used at the time (1994 Q4), we find at horizons of one, two, and four quarters that forecast accuracy increases when we forecast real disposable income in the VAR but not as much as Ireland's forecasts of real labor income.¹³ At an eight-quarter horizon, we do not show forecast improvement, similar to Ireland. We take this as evidence that disposable personal income is a reasonable, albeit noisier, stand-in for labor income. In a preview of our results to follow, we find no forecast improvement from the personal saving rate in real time (RT and RTLTV). Indeed, performance worsens when the saving rate is added.

Forecasts with levels and first differences. We now proceed to our main forecast comparisons for disposable income in Tables 5 and 6. The data used for the estimations begin in 1959 Q1 and our first forecast begins with 1971 Q1. The test used is the ratio of the root-mean-square error (RMSE) of the out-of-sample bivariate VAR forecast to the RMSE of the out-of-sample univariate AR forecast. This is performed forecasting one, two, four, and eight quarters ahead, with tests taken separately at each horizon. We use six lags and lag lengths chosen using the Akaike information criterion (AIC) and the Schwarz information criterion (SIC). In Table 5, we show the regression results using the level of the saving rate, while, in Table 6, we use first differences.

VARs in levels. The first three rows of Table 5 show that in our full sample, 44 quarters longer than Ireland's, the VAR including the level of the personal saving rate outperforms the univariate equation using latest vintage data in four cases out of 12. Using the AIC, the level of the saving rate adds information to the forecast one step ahead, two steps ahead, and four steps

¹³ Ireland gave ratios of MSE, so we have taken square roots. These data refer to unconstrained forecasts.

ahead. In real time, by contrast, there is essentially no advantage to using the level of the personal saving rate. The only forecast improvement is in using the AIC in one-step-ahead forecasts, and the forecast improvement is only 0.4 percent.

Two factors undermining forecast accuracy are the downward trend in the saving rate after 1981 and that the advance estimates of the saving rate are only weakly correlated with the final saving rates. We can eliminate both factors by focusing on latest vintage data before 1982, when the saving rate begins its long-term decline.

We find that there is value to using the level of the saving rate in the first part of the sample, 1971 to 1981, using latest vintage data, with 10 out of 12 forecasts showing improvement. Forecast improvements using the AIC are quite large for the one-, two-, and four-step-ahead forecasts. A key observation is that forecast improvement disappears at longer horizons when we use real-time data. However, in the case of one-step-ahead forecasts, using the AIC or the SIC, there is forecast improvement even with real-time data. The asterisks appearing in Table 5 show cases in which the VAR using the level of the saving rate outperforms the VAR using the first difference. When the data are sufficiently revised, such as the latest vintage data for the period 1971 to 1981, the level of the saving rate is more informative than the first difference: The model using the level outperforms the model using the first difference in 10 out of 12 cases, and in all eight cases when the AIC or the SIC is used to choose lag lengths. In sharp contrast, in real time, the first difference regression is almost always superior.

From 1982 forward, while the latest vintage of data suggests value to including the level of the personal saving rate for one-step-ahead forecasts (for six lags and AIC chosen lags), the improvement is modest and completely disappears in real time.

First difference results. The first three rows of Table 6 show that in our full sample, using the AIC or the SIC, the first difference of the personal saving rate is useful in forecasting real disposable income, at all lag lengths, whether we use latest vintage data or real time data. Only in

one case out of the 24 permutations is there worsening of the forecast. Using six lags, there is improvement only in latest vintage data.

Note further that parsimony is valuable: in Table 6, the AIC and SIC chosen lag lengths almost invariably show improvement, while the six lag VARs show improvement only eight times out of 36. Given that parsimony matters, the first difference regressions perform far better than level regressions. In the full sample, there is only one case out of 36 in which the levels regression outperforms the first difference regression.

Forecasts with PIH restrictions. Let us now return to the VAR estimated on the level of the personal saving rate. Do the data match the PIH theory qualitatively in sample? Consider estimating the equation for the growth rate of real disposable personal income in our VAR on an expanding window of observations, just as we do in our forecasting experiments. Real-time data have a negative sum-of-coefficients for the saving rate for most of the history, providing some evidence that the permanent income hypothesis is correct (Figure 6). However, the sum has become progressively less negative and has actually been near zero since the late 1990s. This suggests either that the empirical validity of this aspect of the permanent income hypothesis has weakened over this period or that the data on saving have become sufficiently noisy that the hypothesis cannot be verified. From a forecasting perspective, noisy measurements of the saving rate could place a premium on imposing some theory-driven restrictions on the data.

Do the PIH restrictions improve forecasts of real disposable personal income? Following Ireland, we investigate whether Campbell's PIH restrictions improve forecasts compared to either the AR or the unrestricted VAR. In the appendix, we demonstrate that the restrictions hold to a close approximation in our VAR estimated on the level and first difference of the saving rate and growth in real income. In Table 7, we see that for latest vintage data, in the pre-1982 period, the PIH-restricted VAR reduces forecast error at all forecast horizons, for all lag lengths. Reductions

are also substantial, ranging from 6 to 16 percent, compared with the previous results shown for the unrestricted VAR in Table 5.¹⁴

In addition, the PIH restricted VAR improves in the pre-1982 sample using real-time data at all lag lengths for one-step-ahead forecasts, as well as in a few other cases. It is noteworthy, however, in the post-1981 sample period, that with PIH restrictions, the VAR does not improve forecasts compared with the univariate autoregressive model for real disposable personal income.

The PIH restrictions are, however, quite beneficial. In Table 8, we compare the PIH restricted VAR forecast performance with the unrestricted VAR. In general, we see that for all time periods and horizons, the PIH restrictions improve forecasts, as Ireland found, particularly for the AIC and SIC selected lag lengths. Notably, we see improvement over the unconstrained VAR, even in real time.

PIH restrictions in first differences. Finally, to complete our analysis, we impose an approximation to the Campbell PIH restrictions on our VAR, estimated on the rate of growth of real disposable income and the first difference of the saving rate, as described in the appendix.

This implies that in our system

$$\begin{pmatrix} \Delta y_t \\ \Delta s_t \end{pmatrix} = \begin{pmatrix} \tilde{a}(L) & \tilde{b}(L) \\ \tilde{c}(L) & \tilde{d}(L) \end{pmatrix} \begin{pmatrix} \Delta y_{t-1} \\ \Delta s_{t-1} \end{pmatrix} + \begin{pmatrix} u_{yt} \\ u_{st} \end{pmatrix},$$

the coefficients across the two equations should be the same, that is,

$$\begin{aligned} \tilde{c}_i &= \tilde{a}_i, & i \geq 1 \\ \tilde{d}_i &= \tilde{b}_i, & i \geq 1. \end{aligned}$$

Imposing these restrictions, as we do in Table 9, we see that forecasts improve over the univariate specification in all but six out of 108 cases. In no cases does the restricted VAR root-mean-square error exceed that of the AR by more than one-half of a percent. This appears to be a very useful methodology for forecasting real disposable

¹⁴ Following Ireland, we set r , the constant real rate of interest, to 0.01 when we impose the PIH restrictions. This implies an annualized rate of 4 percent.

personal income, despite the probability that the level of the savings rate in real time is very noisy.

Conclusion

We have argued that measures of personal saving are subject to substantial measurement error. Benchmark revisions, in particular, have a positive bias. The contention that the low personal saving rate implies that in the future consumption must rise more slowly than income may be wrong: benchmark revisions might well result in the current low rate being revised substantially upward. Taken together, our results suggest that one should be careful about drawing inferences based on the latest observations of the *level* of the U.S. personal saving rate. However, changes in the personal saving rate seem to provide reliable information on future disposable personal income. Imposing restrictions from the permanent income hypothesis also helps to improve the real-time forecasts in VARs estimated without the restrictions.

Appendix

PIH Restrictions in a VAR

Campbell (1987) and Ireland (1995) use the permanent income hypothesis to derive the following equation relating the level of real saving per capita (S) to real labor income per capita (Y_t):

$$S_t = -\sum_{j=1}^{\infty} (1+r)^{-j} E_t \Delta Y_{t,t+j}. \quad (1.1)$$

Campbell, however, uses some tedious algebraic manipulations of (1.1) to derive another equation that is useful for understanding how to impose the cross-equation restrictions on a VAR that are implied by the permanent income hypothesis:

$$S_t - \Delta Y_{t,t} - (1+r)S_{t-1} = -\frac{r}{1+r} \varepsilon_t \quad (1.2)$$

where: $\varepsilon_t = \sum_{j=0}^{\infty} (1+r)^{-j} [E_t Y_{t,t+j} - E_{t-1} Y_{t,t+j}]$.

From the definition of conditional expectations, we have $E_{t-1} \varepsilon_t = 0$. To see the nature of the cross-equation restrictions implied by (1.2), consider the two-equation VAR given by

$$\begin{pmatrix} \Delta Y_{t,t} \\ S_t \end{pmatrix} = \begin{pmatrix} a(L) & b(L) \\ c(L) & d(L) \end{pmatrix} \begin{pmatrix} \Delta Y_{t,t-1} \\ S_{t-1} \end{pmatrix} + \begin{pmatrix} u_{Yt} \\ u_{St} \end{pmatrix}. \quad (1.3)$$

The terms $a(L)$, $b(L)$, $c(L)$, and $d(L)$ are coefficient polynomials in the lag operator of the form $x(L) = \sum_{j=1}^p x_j L^{j-1}$. The PIH places restrictions on these coefficients.

Taking conditional expectations in (1.2), using $E_{t-1} \varepsilon_t = 0$, yields

$$E_{t-1} S_t - E_{t-1} \Delta Y_{t,t} - (1+r)S_{t-1} = 0 \quad (1.4)$$

where $E_{t-1} S_{t-1} = S_{t-1}$. Now, we can use the VAR to form expressions for $E_{t-1} S_t$ and $E_{t-1} \Delta Y_{t,t}$ in (1.4), yielding

$$\begin{aligned}
& \left\{ \left[c_1 \Delta Y_{l,t-1} + \dots + c_p \Delta Y_{l,t-p} \right] + \left[d_1 S_{t-1} + \dots + d_p S_{t-p} \right] \right\} - \\
& \left\{ \left[a_1 \Delta Y_{l,t-1} + \dots + a_p \Delta Y_{l,t-p} \right] + \left[b_1 S_{t-1} + \dots + b_p S_{t-p} \right] \right\} - \\
& (1+r) S_{t-1} = 0.
\end{aligned} \tag{1.5}$$

Expression (1.5) must hold for all time periods, implying the following set of restrictions:

$$\begin{aligned}
c_j &= a_j, \quad j \geq 1 \\
d_1 &= b_1 + (1+r) \\
d_j &= b_j, \quad j \geq 2.
\end{aligned} \tag{1.6}$$

We use a variant of the VAR (1.3) given by

$$\begin{pmatrix} \Delta y_t \\ s_t \end{pmatrix} = \begin{pmatrix} a(L) & b(L) \\ c(L) & d(L) \end{pmatrix} \begin{pmatrix} \Delta y_{t-1} \\ s_{t-1} \end{pmatrix} + \begin{pmatrix} u_{yt} \\ u_{st} \end{pmatrix}, \tag{1.7}$$

where Δy_t is the first difference of the log of real disposable personal income and s_t (lower case) is the personal saving *rate*. Thus, our VAR makes the following substitutions: (1) Our measure of income is real disposable income, not real labor income; (2) Our measure of income enters as a growth rate (log first difference), not as the first difference of a level; and (3) we use the saving rate (s), not the level of real saving per capita (S). In what follows, we ignore the difference in the income concept and focus instead on the difference between income growth (Δy) and the first difference of the level (ΔY). We also focus on our use of the saving rate, not the level of saving.

Can we still impose PIH restrictions on our VAR (1.7)? Perhaps so, as the following argument suggests. Consider dividing both sides of (1.2) by Y_{t-1} , to yield:

$$\frac{S_t}{Y_{t-1}} - \frac{\Delta Y_t}{Y_{t-1}} - (1+r) \frac{S_{t-1}}{Y_{t-1}} = -\frac{r}{1+r} \frac{\varepsilon_t}{Y_{t-1}}. \tag{1.8}$$

Note that $S_t/Y_{t-1} \approx s_t$, $\Delta Y_t/Y_{t-1} \approx \Delta y_t$, $S_{t-1}/Y_{t-1} \equiv s_{t-1}$, and that $E_{t-1}(\varepsilon_t/Y_{t-1}) = 0$, so that a reasonable approximation of (1.4) is, in terms of our VAR (1.7), given by

$$E_{t-1} s_t - E_{t-1} \Delta y_t - (1+r) s_{t-1} \approx 0. \tag{1.9}$$

If one is willing to accept (1.9) as a reasonable approximation, the same analysis that led to the restrictions given in (1.6) also applies to our reformulated VAR.

Consider now a variation on our reformulated VAR in which the saving rate enters in first difference form:

$$\begin{pmatrix} \Delta y_t \\ \Delta s_t \end{pmatrix} = \begin{pmatrix} a(L) & b(L) \\ c(L) & d(L) \end{pmatrix} \begin{pmatrix} \Delta y_{t-1} \\ \Delta s_{t-1} \end{pmatrix} + \begin{pmatrix} u_{yt} \\ u_{st} \end{pmatrix}. \quad (1.10)$$

Can we still impose PIH restrictions? Here the case is somewhat harder to make. Using (1.8), we have

$$\frac{S_t}{Y_{t-1}} - \frac{S_{t-1}}{Y_{t-1}} - \frac{\Delta Y_t}{Y_{t-1}} = -\frac{r}{1+r} \frac{\varepsilon_t}{Y_{t-1}} + r \frac{S_{t-1}}{Y_{t-1}}, \quad (1.11)$$

and using our approximations, where we use $\Delta s_t \equiv s_t - s_{t-1}$,

$$\Delta s_t - \Delta y_t \approx -\frac{r}{1+r} \frac{\varepsilon_t}{Y_{t-1}} + r s_{t-1}. \quad (1.12)$$

When the conditional expectation of the right-hand side is small, we have the restrictions that we impose in the paper:

$$\begin{aligned} c_j &= a_j, & j \geq 1 \\ d_j &= b_j, & j \geq 1. \end{aligned} \quad (1.13)$$

References

- Aruoba, Borogan S., "Data Revisions Are Not Well-behaved," Working Paper, December 2005.
- Boskin, Michael J., "Economic Measurement: Progress and Challenges," *American Economic Association Papers and Proceedings*, May 2000, 247-252.
- Bureau of Economic Analysis, U.S. Department of Commerce, *Survey of Current Business*, various issues, 1965 to 2005.
- Campbell, John Y., "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Personal Income Hypothesis," *Econometrica* 55, November 1987, 1249-73.
- Cole, Rosanne, "Data Errors and Forecasting Accuracy," in Jacob Mincer, ed., *Economic Forecasts and Expectations: Analyses of Forecasting Behavior and Performance*, New York: NBER, 1969, 47-82.
- Croushore, Dean, "Forecasting with Real-Time Macroeconomic Data," in Graham Elliott, Clive W.J. Granger, and Allen Timmermann, eds., *Handbook of Economic Forecasting*, Amsterdam: North-Holland, 2006.
- Croushore, Dean, and Tom Stark, "A Real-Time Data Set for Macroeconomists: Does the Data Vintage Matter?," *Review of Economics and Statistics* 85, August 2003, 605-617.
- Croushore, Dean and Tom Stark, "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics* 105, November 2001, 111-130.
- Diebold, Francis X., and Glenn D. Rudebusch, "Forecasting Output with the Composite Leading Index: A Real-Time Analysis," *Journal of the American Statistical Association* 86, September 1991, 373-393.
- Gale, William G., and John Sabelhaus, "Perspectives on the Household Saving Rate," *Brookings Papers on Economic Activity* 1, 1999, 181-224.
- Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy* 86, December 1978, 971-987.

- Howrey, E. Philip, "The Use of Preliminary Data in Econometric Forecasting," *Review of Economics and Statistics* 60, May 1978, 193-200.
- Ireland, Peter N., "Using the Permanent Income Hypothesis for Forecasting," *Federal Reserve Bank of Richmond Economic Quarterly* 81, Winter 1995, 49-63.
- King, Robert G., "Quantitative Theory and Econometrics," manuscript, University of Virginia, 1995.
- Koenig, Evan, Sheila Domas, and Jeremy Piger, "The Use and Abuse of Real-Time Data in Economic Forecasting," *Review of Economics and Statistics*, 85 2003, 618-628.
- Mankiw, N. Gregory, and Matthew D. Shapiro, "News or Noise: An Analysis of GNP Revisions," *Survey of Current Business* 66, May 1986, 20-25.
- Stark, Tom, and Dean Croushore, "Forecasting with a Real-Time Data Set for Macroeconomicists," *Journal of Macroeconomics* 24, December 2002, 507-31.
- Taubman, Paul, "Personal Saving: A Time Series Analysis of Three Measures of the Same Conceptual Series," *Review of Economics and Statistics* 50, February 1968, 125-129.

Table 1. Personal Saving Rate, Five-Year Averages, After Benchmark Revisions, Percentage Points

Five-year periods	Advance Estimate	Vintage						
		76Q1	81Q1	86Q1	93Q1	97Q2	00Q2	04Q1
65 Q3 to 70 Q2	6.30	6.57	7.21*	7.15	7.20	7.83*	8.55*	8.58
70 Q3 to 75 Q2	7.32	7.53	8.08*	8.71*	8.40	8.94*	10.09*	10.10
75 Q3 to 80 Q2	5.59		5.98	7.20*	7.10	7.68*	9.27*	9.14
80 Q3 to 85 Q2	5.49			6.52	7.98*	8.48*	10.25*	10.37
85 Q3 to 90 Q2	4.33				4.76	5.67*	7.80*	7.45
90 Q3 to 95 Q2	4.34					5.14	7.41*	6.32
95 Q3 to 00 Q2	2.69							3.53
00 Q3 to 05 Q2	1.78							

*More than 0.5 percentage point larger than in the previous benchmark revision.

Source: BEA, Federal Reserve Bank of Philadelphia Real-Time Data Set for Macroeconomists

Table 2A. Regression coefficients for: $PSR_{V,t} = \alpha + \beta PSR_{A,t} + e_t$

(Standard errors in parentheses are Newey-West HAC standard errors, lag truncation=3)

Time Period		Vintage				
		1985q3	1990q3	1995q3	2000q3	2005q3
1965Q3 to 1985Q2	α	2.28* (0.54)	3.26* (0.71)	5.10* (0.64)	7.79* (0.78)	7.56* (0.70)
	β	0.747** (0.08)	0.665** (0.11)	0.416** (0.09)	0.284** (0.12)	0.321** (0.10)
	R^2	0.57	0.45	0.25	0.10	0.12
1970Q3 to 1990Q2	α		0.76 (0.88)	2.94* (0.82)	6.43* (0.62)	5.97* (0.65)
	β		1.026 (0.13)	0.724** (0.12)	0.514** (0.09)	0.579** (0.10)
	R^2		0.62	0.40	0.36	0.36
1975Q3 to 1995Q2	α			1.70* (0.87)	5.07* (0.85)	3.72* (1.12)
	β			0.893 (0.16)	0.732 (0.16)	0.932 (0.21)
	R^2			0.35	0.33	0.33
1980Q3 to 2000Q2	α				2.33* (0.86)	2.45* (0.72)
	β				1.173 (0.19)	1.059 (0.19)
	R^2				0.53	0.43
1985Q3 to 2005Q2	α					2.07* (0.55)
	β					0.830 (0.15)
	R^2					0.41

*greater than 0, p value < .01

**less than 1, p value < .01

***less than 1, p value < .05

Table 2B. Regression coefficients for: $DPSR_{Y,t} = \alpha + \beta DPSR_{A,t} + e_t$
 (Standard errors in parentheses are Newey-West HAC standard errors, lag truncation=3)

Time Period		Vintage				
		1985q3	1990q3	1995q3	2000q3	2005q3
1965Q3 to 1985Q2	α	0.025 (0.044)	0.030 (0.047)	0.056 (0.048)	0.066 (0.048)	0.066 (0.051)
	β	0.809** (0.072)	0.757** (0.060)	0.787** (0.063)	0.732** (0.067)	0.762** (0.061)
	R^2	0.73	0.67	0.67	0.66	0.65
1970Q3 to 1990Q2	α		0.018 (0.049)	0.011 (0.056)	0.030 (0.052)	0.023 (0.055)
	β		0.794** (0.062)	0.796** (0.064)	0.700** (0.067)	0.735** (0.065)
	R^2		0.69	0.67	0.64	0.65
1975Q3 to 1995Q2	α			0.000 (0.054)	-0.019 (0.053)	-0.021 (0.055)
	β			0.703** (0.087)	0.626** (0.085)	0.682** (0.085)
	R^2			0.49	0.45	0.46
1980Q3 to 2000Q2	α				-0.068 (0.052)	-0.037 (0.057)
	β				0.679** (0.082)	0.667** (0.100)
	R^2				0.46	0.39
1985Q3 to 2005Q2	α					-0.066 (0.052)
	β					0.754** (0.090)
	R^2					0.54

Table 3
Forecasts of Real Personal Consumption Expenditure Growth Using the Level of the Personal Saving Rate

Model 1:
$$\begin{pmatrix} \Delta c_t \\ s_t \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} + \begin{pmatrix} \phi_{11}(L) & \phi_{12}(L) \\ \phi_{21}(L) & \phi_{22}(L) \end{pmatrix} \begin{pmatrix} \Delta c_{t-1} \\ s_{t-1} \end{pmatrix} + \begin{pmatrix} u_{ct} \\ u_{st} \end{pmatrix}$$

Model 2: $\Delta c_t = \mu + \phi(L)\Delta c_{t-1} + u_t$

Data from 1959Q1 to 2005Q2

Forecast variable: Percent changes of real personal consumption at annual rates

RMSE (model 1)/RMSE (model 2)

Numbers below unity (bold italics) mean the saving rate improves the forecast for c.

	1 Step Ahead			2 Steps Ahead			4 Steps Ahead			8 Steps Ahead		
	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC
1971:Q1 – 2005:Q2 (Full Sample)												
RT	1.050	1.008	1.022	1.051	1.034	1.044	1.112	1.059	1.054	1.092	1.085	1.038
LV	1.125	1.088	1.050	1.127	1.099	1.085	1.159	1.091	1.098	1.165	1.091	1.076
RTL	1.063	1.016	1.029	1.061	1.033	1.042	1.122	1.049	1.046	1.160	1.094	1.051
1971:Q1 – 1981:Q4 (Pre-1982)												
RT	1.062	0.993	1.004	1.048	1.025	1.030	1.109	1.029	1.031	1.054	1.052	1.006
LV	1.164	1.134	1.073	1.150	1.134	1.117	1.195	1.112	1.127	1.203	1.104	1.088
RTL	1.064	1.012	1.003	1.025	1.028	1.012	1.087	1.032	1.001	1.050	1.058	0.997
1982:Q1 – 2005:Q2 (Post-1981)												
RT	1.023	1.037	1.054	1.061	1.057	1.077	1.119	1.121	1.098	1.153	1.139	1.088
LV	1.047	0.999	1.010	1.069	1.022	1.021	1.095	1.054	1.048	1.106	1.070	1.059
RTL	1.061	1.024	1.074	1.140	1.042	1.096	1.181	1.079	1.114	1.313	1.147	1.124

Each entry is the ratio of out-of-sample RMSE of the VAR model (with lags of the personal saving rate) to the RMSE of univariate AR model. *Numbers below unity (bold italics) mean the saving rate improves the forecast for consumption growth.* LV indicates that the latest vintage of data (2005 Q3) was used to estimate and forecast the model and to compute the forecast errors. RTL indicates that real-time data were used to estimate and forecast the model, but the latest vintage of data was used to compute forecast errors. RT indicates that real-time data were used to estimate and forecast the model and to evaluate forecast errors. “One step ahead” refers to forecasts for the one-step-ahead quarter-over-quarter percent change. “Two steps ahead” refers to forecasts for the two-step-ahead two-quarter average percent change. “Four steps ahead” and “eight steps ahead” are defined in a similar manner.

Table 4. Comparing Ireland's forecasts of the change in real labor income per capita with our forecasts of the percent change in real disposable income

Numbers below unity (bold italics) mean the saving rate improves the forecast for income.

Data from 1959 Q1 to 1994 Q3

Forecast period: 1971 Q1 to 1994 Q3

Forecast period	Forecast horizon			
Forecast period: 1971 Q1 to 1994 Q3	1 Quarter Ahead	2 Quarters Ahead	4 Quarters Ahead	8 Quarters Ahead
Forecast variable: Total change in real labor income per capita (Source: Ireland, 1995)				
LV	.97	.95	.90	1.07
Forecast variable: Percent change in real disposable income				
RT	1.050	1.068	1.078	1.113
LV	.974	.982	.956	1.065
RTL	1.047	1.072	1.094	1.158

Each entry is the ratio of out-of-sample RMSE of the VAR model (with lags of the personal saving rate) to the RMSE of the univariate AR model, using six lags.

LV indicates that the latest vintage of data (2004 Q4) available at the time of Ireland's study was used to estimate and forecast the model and to compute the forecast errors. RTL indicates that real-time data were used to estimate and forecast the model, but the latest vintage of data was used to compute forecast errors. RT indicates that real-time data were used to estimate and forecast the model and to evaluate forecast errors.

Ireland forecasted the total change in real labor income per capita. We forecast the percent change in real disposable income. "One quarter ahead" refers to forecasts for the one-step-ahead quarter-over-quarter change or percent change. "Two quarters ahead" refers to forecasts for the two-step-ahead two-quarter average change or percent change. "Four quarters ahead" and "eight quarters ahead" are defined in a similar manner.

Table 5
Forecasts of Real Disposable Personal Income Growth Using Level of Personal Saving Rate

Model 1:
$$\begin{pmatrix} \Delta y_t \\ s_t \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} + \begin{pmatrix} \phi_{11}(L) & \phi_{12}(L) \\ \phi_{21}(L) & \phi_{22}(L) \end{pmatrix} \begin{pmatrix} \Delta y_{t-1} \\ s_{t-1} \end{pmatrix} + \begin{pmatrix} u_{yt} \\ u_{st} \end{pmatrix}$$

Model 2: $\Delta y_t = \mu + \phi(L)\Delta y_{t-1} + u_t$

Data from 1959Q1 to 2005Q2

Forecast variable: Percent changes of real disposable income at annual rates

RMSE (model 1)/RMSE (model 2)
Numbers below unity (bold italics) mean the saving rate improves the forecast for y.

	1 Step Ahead			2 Steps Ahead			4 Steps Ahead			8 Steps Ahead		
	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC
1971:Q1 – 2005:Q2 (Full Sample)												
RT	1.029	<i>0.996</i>	1.018	1.069	1.092	1.133	1.086	1.124	1.190	1.119	1.202	1.297
LV	<i>0.978*</i>	<i>0.928</i>	1.013	1.010	<i>0.968</i>	1.061	1.015	<i>0.975</i>	1.095	1.102	1.118	1.206
RTL	1.010	<i>0.996</i>	1.007	1.052	1.069	1.110	1.100	1.117	1.183	1.205	1.244	1.350
1971:Q1 – 1981:Q4 (Pre-1982)												
RT	1.024*	<i>0.963</i>	<i>0.947</i>	1.049	1.050	1.056	1.066	1.049	1.075	1.147	1.115	1.131
LV	<i>0.968*</i>	<i>0.882*</i>	<i>0.959*</i>	<i>0.989*</i>	<i>0.919*</i>	<i>0.971*</i>	1.005	<i>0.889*</i>	<i>0.941*</i>	1.108	<i>0.977*</i>	<i>0.967*</i>
RTL	1.002*	<i>0.967*</i>	<i>0.951</i>	1.037	1.043	1.044	1.110	1.071	1.079	1.267	1.167	1.175
1982:Q1 – 2005:Q2 (Post-1981)												
RT	1.036	1.040	1.111	1.104	1.150	1.240	1.126	1.245	1.364	1.073	1.317	1.501
LV	<i>0.989</i>	<i>0.978</i>	1.064	1.039	1.025	1.156	1.029	1.060	1.234	1.095	1.237	1.402
RTL	1.023	1.029	1.075	1.076	1.103	1.195	1.083	1.171	1.299	1.123	1.323	1.519

Each entry is the ratio of out-of-sample RMSE of the VAR model (with lags of the personal saving rate) to the RMSE of the univariate AR model

* indicates a lower ratio than corresponding forecasts with first differences (Table 6). *Numbers below unity (bold italics) mean the saving rate improves the forecast for y.* LV indicates that the latest vintage of data (2005Q3) was used to estimate and forecast the model and to compute the forecast errors. RTL indicates that real-time data were used to estimate and forecast the model, but the latest vintage of data was used to compute forecast errors. RT indicates that real-time data were used to estimate and forecast the model and to evaluate forecast errors.

Table 6
Forecasts of Real Disposable Personal Income Growth Using First Difference of Personal Saving Rate

Model 1:
$$\begin{pmatrix} \Delta y_t \\ \Delta s_t \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} + \begin{pmatrix} \phi_{11}(L) & \phi_{12}(L) \\ \phi_{21}(L) & \phi_{22}(L) \end{pmatrix} \begin{pmatrix} \Delta y_{t-1} \\ \Delta s_{t-1} \end{pmatrix} + \begin{pmatrix} u_{yt} \\ u_{st} \end{pmatrix}$$

Model 2: $\Delta y_t = \mu + \phi(L)\Delta y_{t-1} + u_t$

Data from 1959Q1 to 2005Q2

Forecast variable: Percent changes of real disposable income at annual rates

RMSE (model 1)/RMSE (model 2)
Numbers below unity (bold italics) mean the saving rate improves the forecast for y.

	1 Step Ahead			2 Steps Ahead			4 Steps Ahead			8 Steps Ahead		
	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC
1971:Q1 – 2005:Q2 (Full Sample)												
RT	1.019	<i>0.947</i>	<i>0.922</i>	1.035	<i>0.986</i>	<i>0.985</i>	1.012	<i>0.954</i>	<i>0.967</i>	1.042	<i>0.987</i>	<i>0.991</i>
LV	<i>0.980</i>	<i>0.907</i>	<i>0.954</i>	<i>0.985</i>	<i>0.953</i>	<i>0.986</i>	<i>0.955</i>	<i>0.925</i>	<i>0.976</i>	1.002	<i>0.972</i>	<i>0.982</i>
RTL	1.009	<i>0.981</i>	<i>0.947</i>	1.023	1.007	<i>0.997</i>	1.015	<i>0.956</i>	<i>0.965</i>	1.064	<i>0.981</i>	<i>0.989</i>
1971:Q1 – 1981:Q4 (Pre-1982)												
RT	1.030	<i>0.950</i>	<i>0.894</i>	1.028	<i>0.994</i>	<i>0.989</i>	1.003	<i>0.936</i>	<i>0.966</i>	1.062	<i>0.991</i>	<i>0.998</i>
LV	1.011	<i>0.902</i>	<i>0.974</i>	1.011	<i>0.961</i>	1.010	<i>0.993</i>	<i>0.913</i>	<i>0.991</i>	1.030	<i>0.988</i>	<i>0.992</i>
RTL	1.010	<i>0.986</i>	<i>0.922</i>	1.000	1.012	1.000	1.008	<i>0.931</i>	<i>0.951</i>	1.095	<i>0.976</i>	<i>0.995</i>
1982:Q1 – 2005:Q2 (Post-1981)												
RT	1.000	<i>0.943</i>	<i>0.961</i>	1.048	<i>0.974</i>	<i>0.979</i>	1.031	<i>0.985</i>	<i>0.971</i>	1.009	<i>0.982</i>	<i>0.981</i>
LV	<i>0.941</i>	<i>0.912</i>	<i>0.933</i>	<i>0.948</i>	<i>0.944</i>	<i>0.958</i>	<i>0.903</i>	<i>0.939</i>	<i>0.960</i>	<i>0.971</i>	<i>0.956</i>	<i>0.971</i>
RTL	1.008	<i>0.974</i>	<i>0.979</i>	1.058	1.000	<i>0.994</i>	1.026	<i>0.987</i>	<i>0.982</i>	1.024	<i>0.986</i>	<i>0.983</i>

Each entry is the ratio of out-of-sample RMSE of the VAR model (with lags of the personal saving rate) to the RMSE of univariate AR model. *Numbers below unity (bold italics) mean the saving rate improves the forecast for y.* LV indicates that the latest vintage of data (2005Q3) was used to estimate and forecast the model and to compute the forecast errors. RTL indicates that real-time data were used to estimate and forecast the model, but the latest vintage of data was used to compute forecast errors. RT indicates that real-time data were used to estimate and forecast the model and to evaluate forecast errors.

Table 7
Forecasts of Real Disposable Personal Income Growth Adding PIH Restrictions Using Levels of Personal Saving Rate

Model 1:
$$\begin{pmatrix} \Delta y_t \\ s_t \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} + \begin{pmatrix} \tilde{a}(L) & \tilde{b}(L) \\ \tilde{c}(L) & \tilde{d}(L) \end{pmatrix} \begin{pmatrix} \Delta y_{t-1} \\ s_{t-1} \end{pmatrix} + \begin{pmatrix} u_{yt} \\ u_{st} \end{pmatrix}$$
 $\tilde{a}_i = \tilde{c}_i, i \geq 1; \quad \tilde{d}_1 = \tilde{b}_1 + (1 + 0.01); \quad \tilde{d}_i = \tilde{b}_i, i \geq 2$

Model 2: $\Delta y_t = \mu + \phi(L)\Delta y_{t-1} + u_t$

Data from 1959Q1 to 2005Q2

Forecast variable: Percent changes of real disposable income at annual rates

RMSE (model 1)/RMSE (model 2)

Numbers below unity (bold italics) mean the saving rate improves the forecast for y.

	1 Step Ahead			2 Steps Ahead			4 Steps Ahead			8 Steps Ahead		
	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC
1971:Q1 – 2005:Q2 (Full Sample)												
RT	<i>0.992</i>	<i>0.981</i>	<i>0.992</i>	1.050	1.063	1.083	1.052	1.090	1.129	1.096	1.138	1.193
LV	<i>0.980</i>	<i>0.928</i>	<i>0.990</i>	1.027	<i>0.981</i>	1.030	1.008	<i>0.998</i>	1.061	1.072	1.109	1.133
RTL	<i>0.978</i>	<i>0.974</i>	<i>0.973</i>	1.033	1.032	1.052	1.037	1.065	1.103	1.101	1.142	1.205
1971:Q1 – 1981:Q4 (Pre-1982)												
RT	<i>0.965*</i>	<i>0.954</i>	<i>0.924</i>	1.015	1.020	1.017	<i>0.994</i>	1.025	1.043	1.021	1.043	1.062
LV	<i>0.928*</i>	<i>0.840*</i>	<i>0.912*</i>	<i>0.935*</i>	<i>0.851*</i>	<i>0.897*</i>	<i>0.878*</i>	<i>0.839*</i>	<i>0.887*</i>	<i>0.912*</i>	<i>0.909*</i>	<i>0.900*</i>
RTL	<i>0.946*</i>	<i>0.943*</i>	<i>0.914</i>	<i>0.997</i>	<i>0.988</i>	<i>0.981</i>	<i>0.997</i>	1.020	1.022	1.049	1.064	1.079
1982:Q1 – 2005:Q2 (Post-1981)												
RT	1.033	1.017	1.081	1.111	1.123	1.175	1.169	1.195	1.262	1.202	1.262	1.356
LV	1.038	1.018	1.063	1.145	1.116	1.164	1.153	1.147	1.214	1.222	1.270	1.323
RTL	1.022	1.010	1.043	1.090	1.088	1.142	1.098	1.116	1.195	1.163	1.221	1.330

Each entry is the ratio of out-of-sample RMSE of the constrained VAR model (with lags of the personal saving rate) to the RMSE of univariate AR model. * indicates a lower ratio than corresponding forecasts with first differences. (Table 9). *Numbers below unity (bold italics) mean the saving rate improves the forecast for y.* LV indicates that the latest vintage of data (2005Q3) was used to estimate and forecast the model and to compute the forecast errors. RTL indicates that real-time data were used to estimate and forecast the model, but the latest vintage of data was used to compute forecast errors. RT indicates that real-time data were used to estimate and forecast the model and to evaluate forecast errors.

Table 8
Forecasts of Real Disposable Personal Income Growth Adding PIH Restrictions Using Levels of Personal Saving Rate

Model 1:
$$\begin{pmatrix} \Delta y_t \\ s_t \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} + \begin{pmatrix} \tilde{a}(L) & \tilde{b}(L) \\ \tilde{c}(L) & \tilde{d}(L) \end{pmatrix} \begin{pmatrix} \Delta y_{t-1} \\ s_{t-1} \end{pmatrix} + \begin{pmatrix} u_{yt} \\ u_{st} \end{pmatrix} \quad \tilde{a}_i = \tilde{c}_i, i \geq 1; \quad \tilde{d}_1 = \tilde{b}_1 + (1 + 0.01); \quad \tilde{d}_i = \tilde{b}_i, i \geq 2$$

Model 2: Model 1 without PIH restrictions.

Data from 1959Q1 to 2005Q2

Forecast variable: Percent changes of real disposable income at annual rates

RMSE (model 1)/RMSE (model 2)

Numbers below unity (bold italics) mean the PIH restrictions improve the forecasts for y relative to the forecasts of the unconstrained VAR.

	1 Step Ahead			2 Steps Ahead			4 Steps Ahead			8 Steps Ahead		
	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC
1971:Q1 – 2005:Q2 (Full Sample)												
RT	<i>0.964</i>	<i>0.985</i>	<i>0.974</i>	<i>0.983</i>	<i>0.974</i>	<i>0.955</i>	<i>0.969</i>	<i>0.970</i>	<i>0.949</i>	<i>0.980</i>	<i>0.947</i>	<i>0.919</i>
LV	1.002	1.000	<i>0.978</i>	1.017	1.013	<i>0.971</i>	<i>0.993</i>	1.024	<i>0.969</i>	<i>0.973</i>	<i>0.992</i>	<i>0.940</i>
RTL	<i>0.968</i>	<i>0.978</i>	<i>0.966</i>	<i>0.982</i>	<i>0.965</i>	<i>0.947</i>	<i>0.943</i>	<i>0.953</i>	<i>0.932</i>	<i>0.914</i>	<i>0.918</i>	<i>0.892</i>
1971:Q1 – 1981:Q4 (Pre-1982)												
RT	<i>0.942</i>	<i>0.991</i>	<i>0.975</i>	<i>0.968</i>	<i>0.971</i>	<i>0.962</i>	<i>0.932</i>	<i>0.978</i>	<i>0.971</i>	<i>0.890</i>	<i>0.935</i>	<i>0.939</i>
LV	<i>0.958</i>	<i>0.953</i>	<i>0.951</i>	<i>0.945</i>	<i>0.926</i>	<i>0.924</i>	<i>0.874</i>	<i>0.944</i>	<i>0.944</i>	<i>0.823</i>	<i>0.931</i>	<i>0.931</i>
RTL	<i>0.944</i>	<i>0.975</i>	<i>0.961</i>	<i>0.962</i>	<i>0.947</i>	<i>0.939</i>	<i>0.898</i>	<i>0.953</i>	<i>0.947</i>	<i>0.829</i>	<i>0.912</i>	<i>0.918</i>
1982:Q1 – 2005:Q2 (Post-1981)												
RT	<i>0.997</i>	<i>0.978</i>	<i>0.973</i>	1.006	<i>0.977</i>	<i>0.948</i>	1.038	<i>0.960</i>	<i>0.926</i>	1.120	<i>0.958</i>	<i>0.903</i>
LV	1.049	1.042	<i>0.999</i>	1.103	1.089	1.007	1.121	1.082	<i>0.984</i>	1.116	1.026	<i>0.944</i>
RTL	1.000	<i>0.981</i>	<i>0.970</i>	1.013	<i>0.986</i>	<i>0.956</i>	1.013	<i>0.953</i>	<i>0.919</i>	1.036	<i>0.923</i>	<i>0.875</i>

Each entry is the ratio of out-of-sample RMSE of the constrained VAR model to the RMSE of the unconstrained VAR model. *Numbers below unity (bold italics) mean the PIH restrictions improve the forecasts for y relative to the forecasts of the unconstrained VAR.* LV indicates that the latest vintage of data (2005Q3) was used to estimate and forecast the model and to compute the forecast errors. RTL indicates that real-time data were used to estimate and forecast the model, but the latest vintage of data was used to compute forecast errors. RT indicates that real-time data were used to estimate and forecast the model and to evaluate forecast errors.

Table 9

Forecasts of Real Disposable Personal Income Growth Adding PIH Restrictions Using First Difference of Personal Saving Rate

$$\text{Model 1: } \begin{pmatrix} \Delta y_t \\ \Delta s_t \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} + \begin{pmatrix} \tilde{a}(L) & \tilde{b}(L) \\ \tilde{c}(L) & \tilde{d}(L) \end{pmatrix} \begin{pmatrix} \Delta y_{t-1} \\ \Delta s_{t-1} \end{pmatrix} + \begin{pmatrix} u_{yt} \\ u_{st} \end{pmatrix} \quad \begin{matrix} \tilde{a}_i = \tilde{c}_i, & i \geq 1 \\ \tilde{d}_i = \tilde{b}_i, & i \geq 1 \end{matrix}$$

$$\text{Model 2: } \Delta y_t = \mu + \phi(L)\Delta y_{t-1} + u_t$$

Data from 1959Q1 to 2005Q2

Forecast variable: Percent changes of real disposable income at annual rates

RMSE (model 1)/RMSE (model 2)

Numbers below unity (bold italics) mean the saving rate improves the forecast for y.

	1 Step Ahead			2 Steps Ahead			4 Steps Ahead			8 Steps Ahead		
	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC	6	AIC	SIC
1971:Q1 – 2005:Q2 (Full Sample)												
RT	<i>0.973</i>	<i>0.940</i>	<i>0.903</i>	1.005	<i>0.989</i>	<i>0.964</i>	<i>0.986</i>	<i>0.975</i>	<i>0.972</i>	<i>0.989</i>	<i>0.990</i>	<i>0.989</i>
LV	<i>0.974</i>	<i>0.909</i>	<i>0.945</i>	<i>0.990</i>	<i>0.963</i>	<i>0.977</i>	<i>0.924</i>	<i>0.960</i>	<i>0.988</i>	<i>0.968</i>	<i>0.985</i>	<i>0.985</i>
RTL	<i>0.973</i>	<i>0.957</i>	<i>0.915</i>	1.001	<i>0.990</i>	<i>0.959</i>	<i>0.952</i>	<i>0.965</i>	<i>0.954</i>	<i>0.982</i>	<i>0.984</i>	<i>0.978</i>
1971:Q1 – 1981:Q4 (Pre-1982)												
RT	<i>0.974</i>	<i>0.944</i>	<i>0.862</i>	1.005	<i>0.999</i>	<i>0.953</i>	<i>0.950</i>	<i>0.964</i>	<i>0.962</i>	1.005	<i>0.990</i>	<i>0.987</i>
LV	<i>0.991</i>	<i>0.883</i>	<i>0.944</i>	<i>0.991</i>	<i>0.939</i>	<i>0.966</i>	<i>0.912</i>	<i>0.936</i>	<i>0.990</i>	<i>0.961</i>	<i>0.987</i>	<i>0.980</i>
RTL	<i>0.964</i>	<i>0.960</i>	<i>0.872</i>	<i>0.989</i>	<i>0.986</i>	<i>0.930</i>	<i>0.929</i>	<i>0.945</i>	<i>0.922</i>	<i>0.987</i>	<i>0.972</i>	<i>0.964</i>
1982:Q1 – 2005:Q2 (Post-1981)												
RT	<i>0.971</i>	<i>0.935</i>	<i>0.958</i>	1.004	<i>0.973</i>	<i>0.981</i>	<i>0.988</i>	<i>0.993</i>	<i>0.989</i>	<i>0.965</i>	<i>0.991</i>	<i>0.992</i>
LV	<i>0.954</i>	<i>0.938</i>	<i>0.945</i>	<i>0.990</i>	<i>0.991</i>	<i>0.989</i>	<i>0.939</i>	<i>0.986</i>	<i>0.986</i>	<i>0.975</i>	<i>0.984</i>	<i>0.990</i>
RTL	<i>0.986</i>	<i>0.954</i>	<i>0.966</i>	1.021	<i>0.996</i>	<i>0.998</i>	<i>0.986</i>	<i>0.989</i>	<i>0.992</i>	<i>0.974</i>	<i>0.996</i>	<i>0.992</i>

Each entry is the ratio of out-of-sample RMSE of the VAR model (with lags of the personal saving rate) to the RMSE of univariate AR model. *Numbers below unity (bold italics) mean the saving rate improves the forecast for y.* LV indicates that the latest vintage of data (2005Q3) was used to estimate and forecast the model and to compute the forecast errors. RTL indicates that real-time data were used to estimate and forecast the model, but the latest vintage of data was used to compute forecast errors. RT indicates that real-time data were used to estimate and forecast the model and to evaluate forecast errors.

Figure 1. Measured Personal Saving Rates

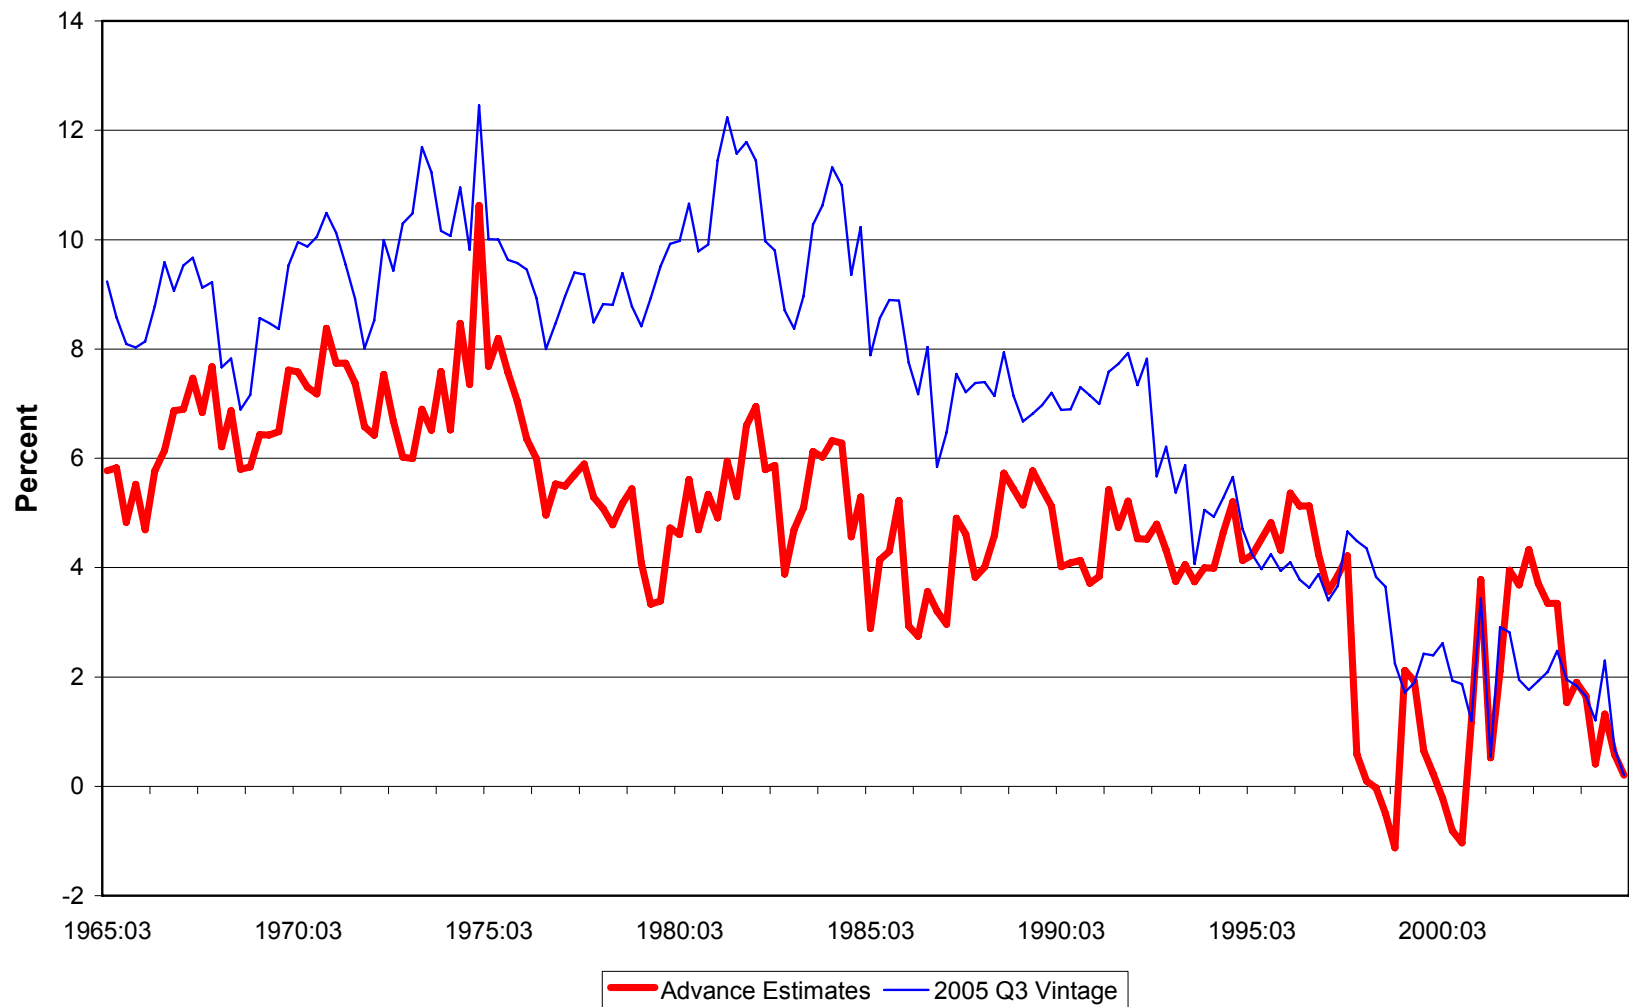


Figure 2
Histogram of Revisions to the Personal Savings Rate
Last-Before-Benchmark minus Advance Estimates

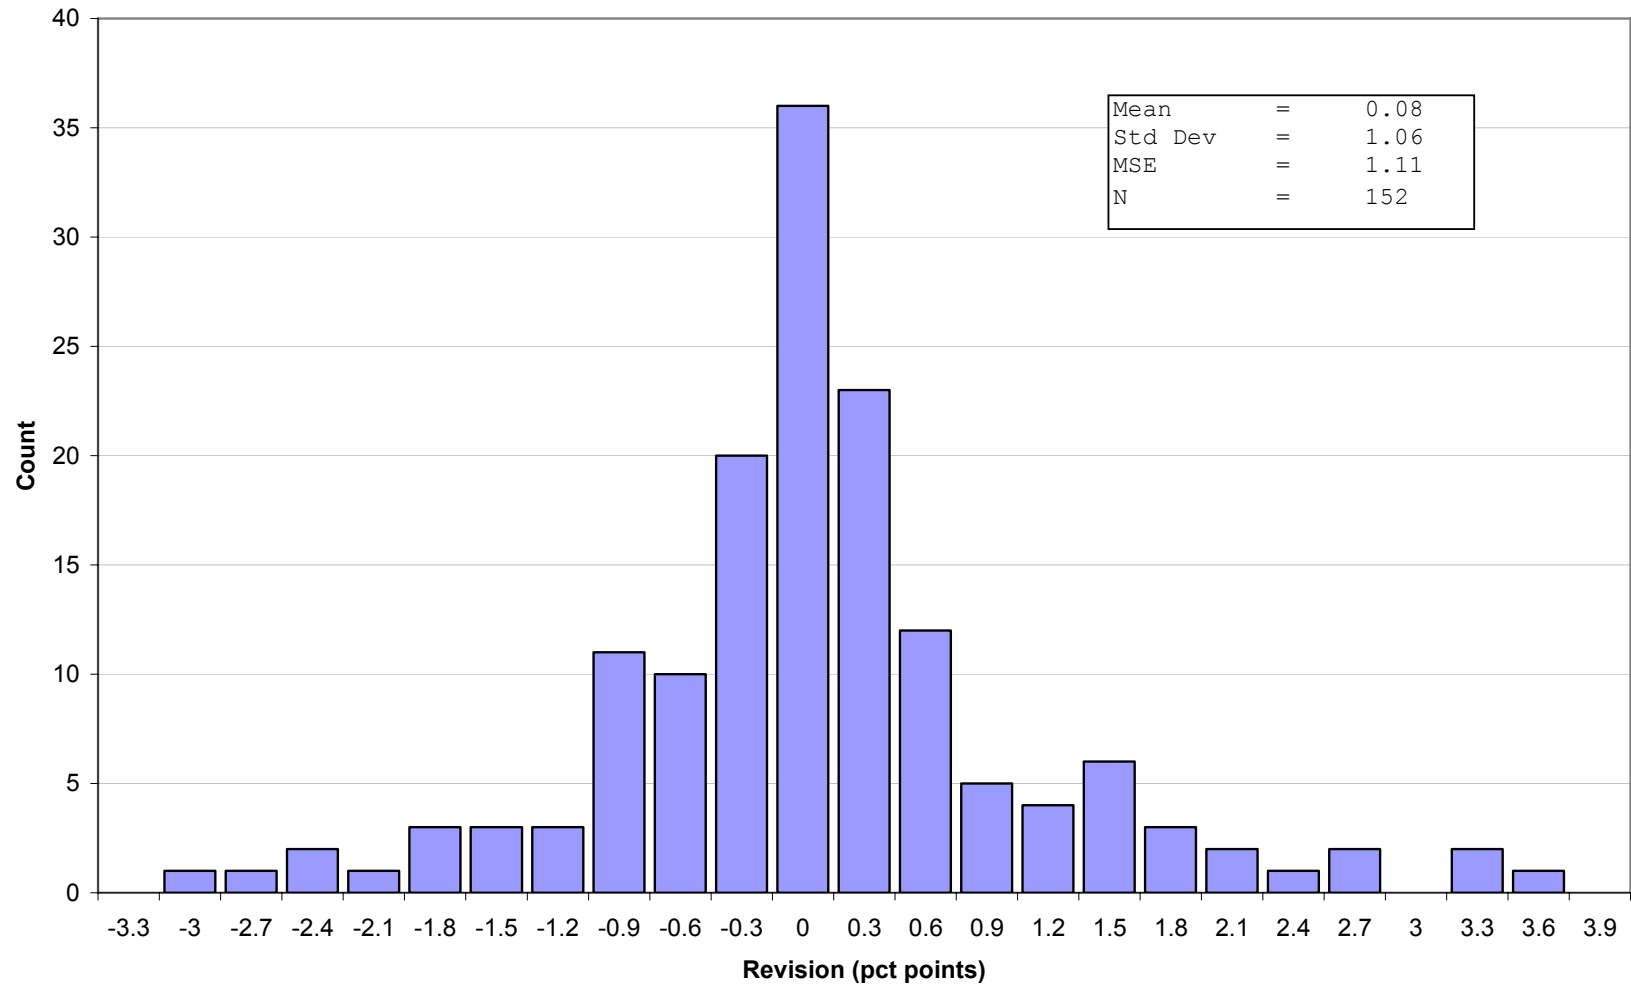


Figure 3
Histogram of Revisions to the Personal Savings Rate
Latest-Available Minus Advance Estimates

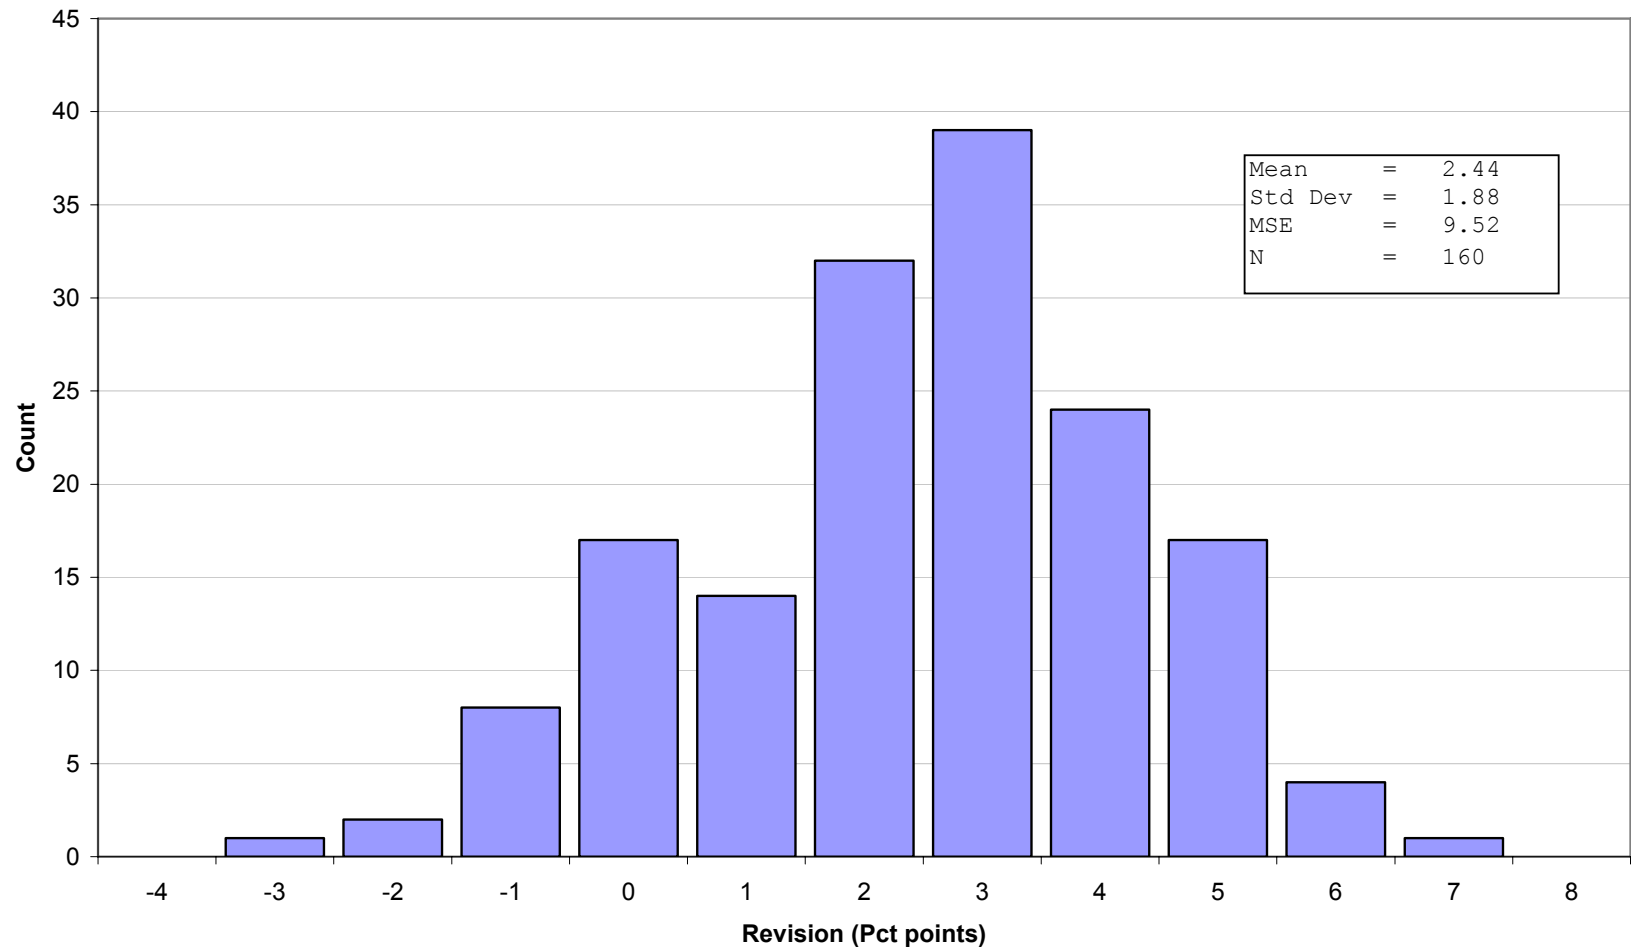


Figure 4
Revisions to Nominal Income and Output

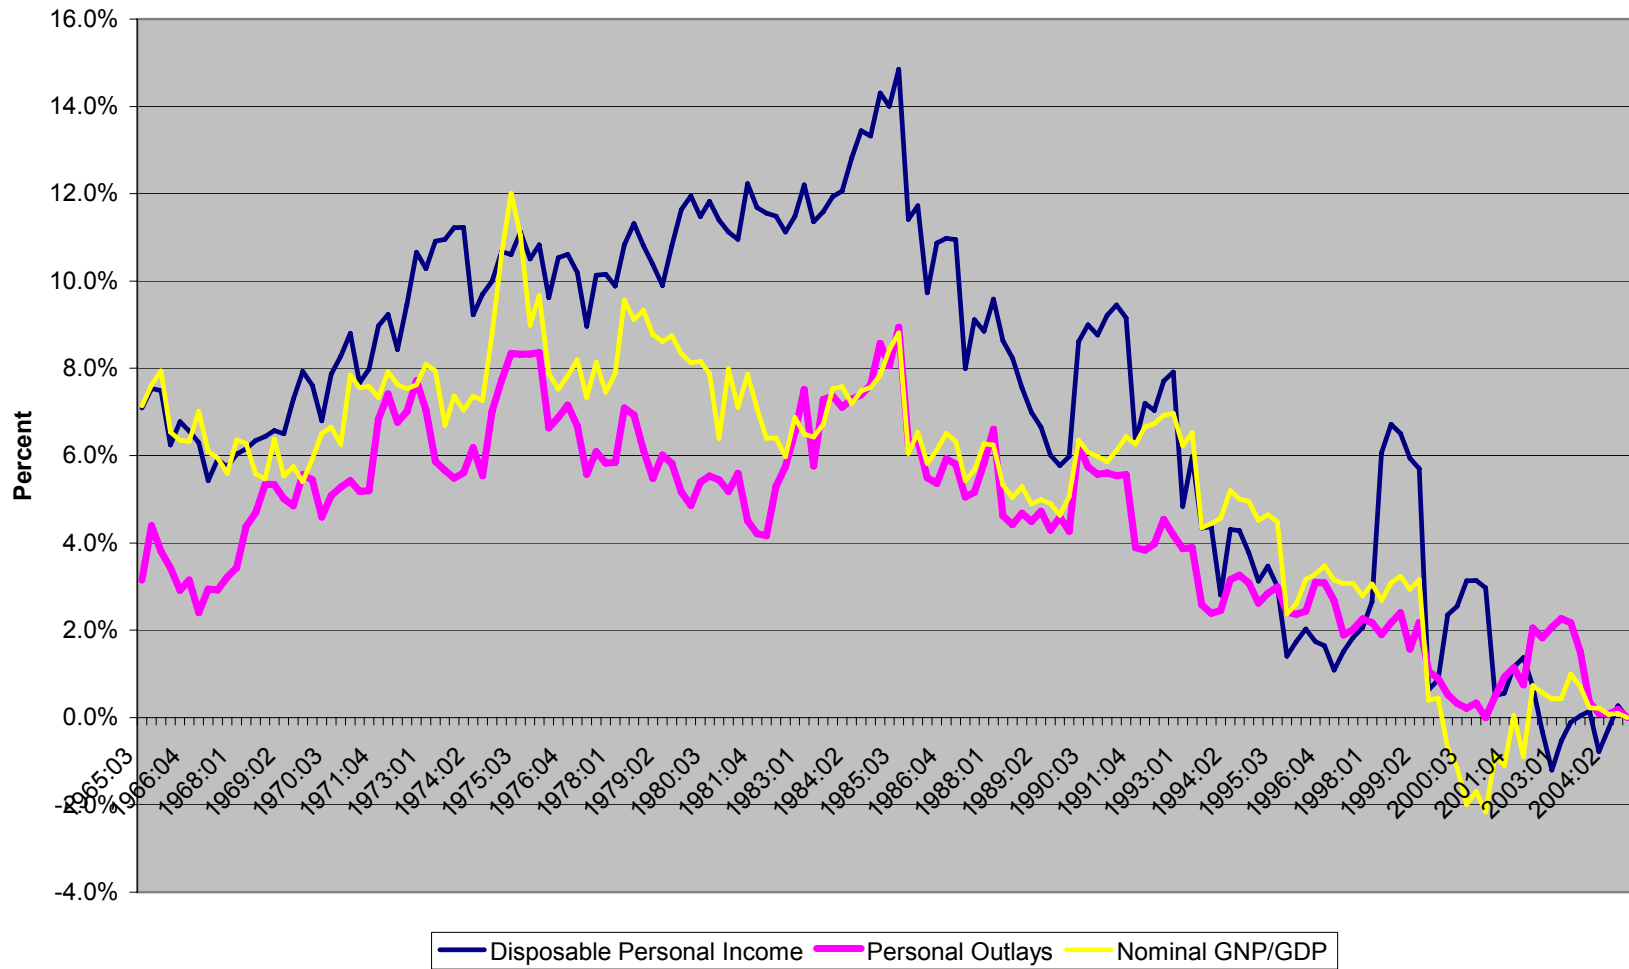
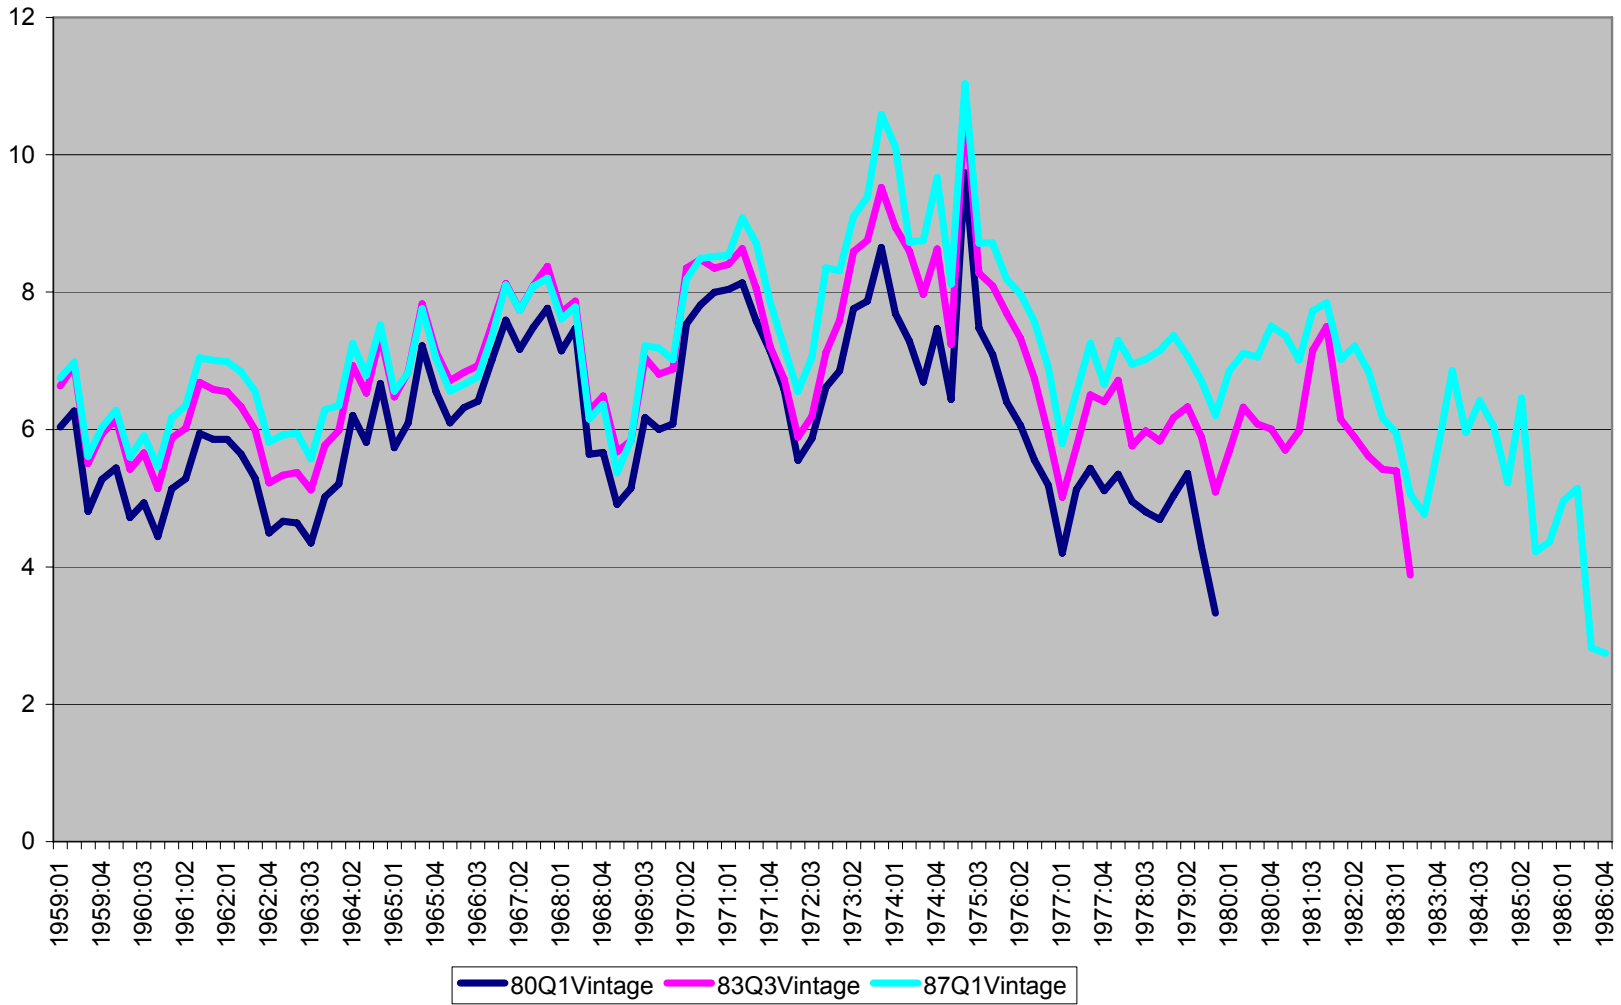


Figure 5a
Recent Saving Rates are Often Below Average



**Figure 5b:
Recent Saving Rates are Often Below Average**

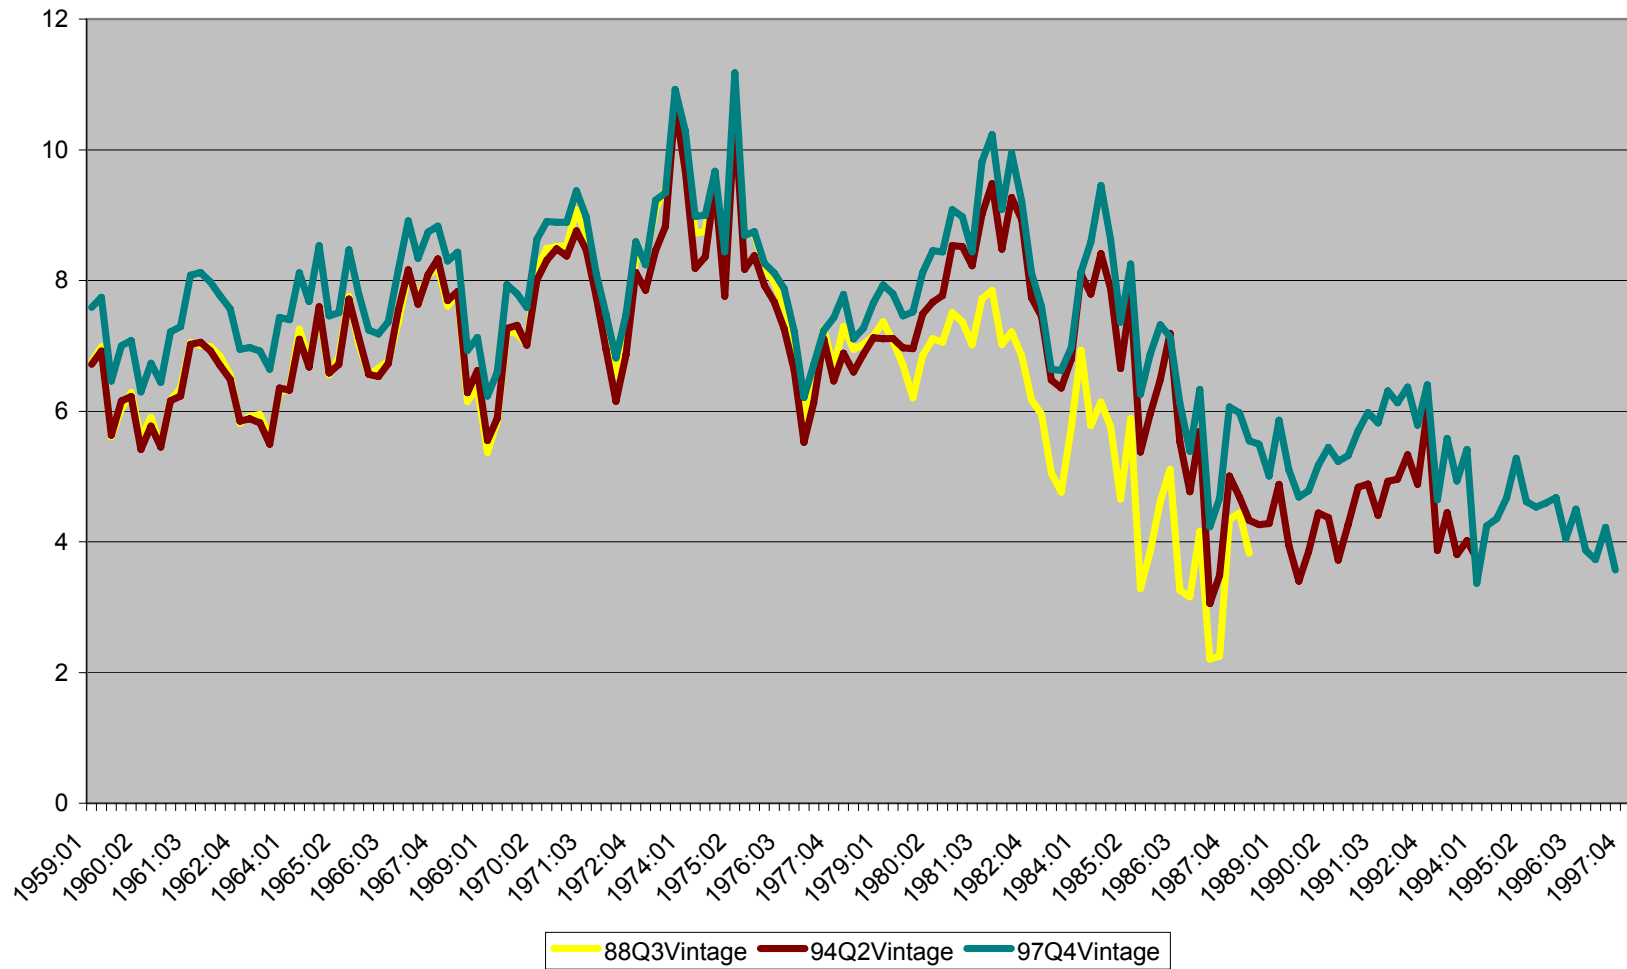


Figure 5c
Recent Saving Rates are Often Below Average

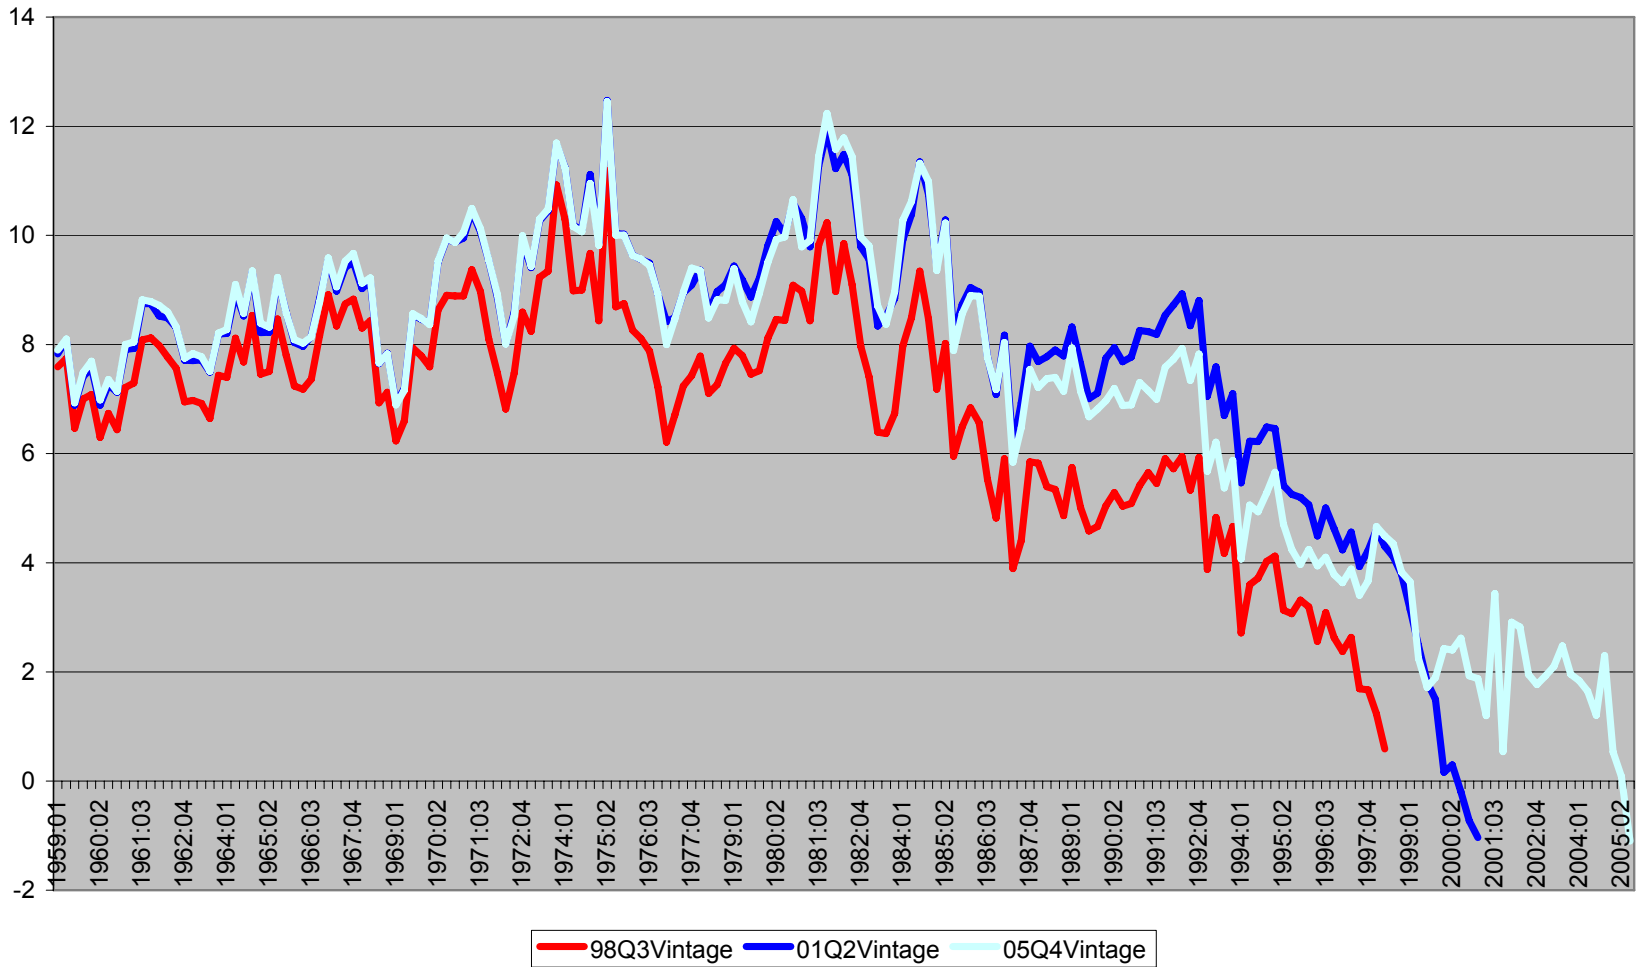
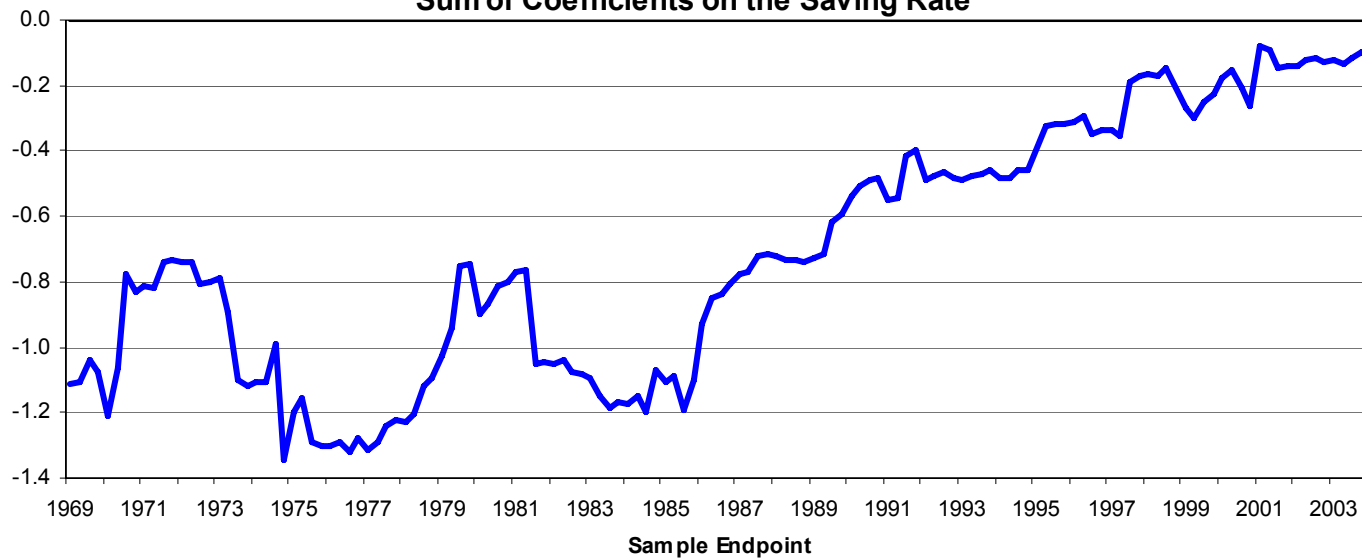


Figure 6.
Sum of Coefficients on the Saving Rate



Notes. The figure plots the sum of coefficients on the saving rate from the VAR equation for real disposable personal income growth. The saving rate is expressed in percentage points. Income growth is expressed in annualized percentage points. Lag length was chosen by the SIC. All estimation begins with the observation for 1959Q1 and adds one additional real-time observation per quarter. The horizontal axis gives the sample endpoints.