The Mismeasured Personal Saving Rate Is Still Useful:
Using Real-Time Data to Improve Forecasting*

BY LEONARD NAKAMURA

People make decisions based on information. Often, with hindsight, they could have made better choices. Economics faces a similar problem: Economic data, when first released, are often inaccurate and may subsequently be revised. In this article, Leonard Nakamura uses the U.S. personal saving rate — a statistic that has often been initially low, then substantially revised upward — to discuss how modern economic statistical techniques can improve forecasting.

People make decisions based on information. A quarterback scanning receivers or a corporate executive concluding a merger must usually make decisions with inadequate information. With hindsight, they often could have made a better choice. A similar problem holds true for economics: Initial economic statistics are often inaccurate and may be subsequently revised as better data become available. One consequence of this process of initial data releases that are later revised is that economists now realize that the quality of economic forecasts needs to be judged against the data available at the time. A second consequence is that when economists make forecasts, they should be aware that the statistics will be revised and incorporate this information into their forecasts.

The U.S. personal saving rate, which has been averaging less than 1 percent of after-tax personal income for the past three years, has often been initially low and then substantially revised upward. I will take this statistic as an example and discuss how modern economic statistical techniques can improve forecasting, by taking into account the difficulties of measuring saving in the short run.

Leonard Nakamura is an assistant vice president and economist in the Philadelphia Fed's Research Department. This article is available free of charge at www.philadelphiafed.org/research-and-data/publications/business-review/.

USING THE SAVING RATE AS A FORECASTING TOOL

Households often make decisions about how much to spend or how much to save based not just on their current income but also on their expectations of future income. To the extent that households base such decisions on expected future income, economists may draw inferences from that behavior and use them to make forecasts about households' income. For example, households may save more when they expect their future income to decline, such as in retirement, and they may save less when they expect their future income to rise. If so, an economist might be able to infer that households expect to retire — and therefore suffer a fall in income — from their saving behavior.

However, understanding the economic behavior in question is only part of the difficulty of forecasting. In practice, economic forecasting suffers from the problem that at the moment a forecast is made, current data on the economy may be imperfect. So the forecaster must try to estimate what will happen tomorrow, not knowing fully what is happening today. As time passes, the data will be improved, but that fact is cold comfort to the forecaster. In the case of data on saving, the personal saving rate is often not measured well initially, making forecasting more difficult.

Averaging across all households in the U.S. economy, we expect household saving to be positive. After all, positive saving supports a rising stock of capital that will make workers more productive. From 1946 to 1992, the personal saving rate was generally stable (Figure 1). If the personal saving rate has been generally stable over time, then whenever the personal saving rate is low, it should tend to rise

*The views expressed here are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

back to its average rate, and vice versa, a process called mean reversion. Since personal saving is defined as after-tax personal income minus personal outlays, a low saving rate seemingly must have one of two implications: Either consumption is expected to fall, or income is expected to rise. The often-expressed view that if saving is low, consumers are overspending and must soon cut back is a tempting one. But as this article will show, the evidence strongly favors the view that if saving is low, it's more likely the case that income is expected to rise.

The saving rate may change for other reasons, as well. For example, households save not only for retirement but also for a rainy day. This "precautionary saving" stems from households' concern that they may suffer a loss of income due to layoffs or ill health. In recent years, financial innovations and the moderation of the severity and frequency of U.S. recessions in recent years may have reduced households' fears of the consequences of income loss. To the extent that changes such as these influence the saving rate, the saving rate will not accurately predict changes in income.

Recently, the measured U.S. saving rate has been very low: under 1 percent of after-tax income in 2005, 2006, and 2007. This stands in contrast to a saving rate of 8.5 percent over the 46 years from 1946 to 1992. Does the fact that the U.S. is experiencing a low saving rate imply swiftly rising income?

It turns out that the current low level of the personal saving rate may well be due to mismeasurement. As I argued in a 2001 *Business Review* article, personal saving is hard to measure and may be understated, particularly over the past 20 years or so, as a result of changes in the way the economy behaves and is measured. In particular, we appear to have systematically undercounted U.S. investment in developing new products, which has resulted in uncounted income and saving. In this situation, is it possible to still use the saving rate to make forecasts?

In our working paper, Tom Stark and I point out that, in the past, initial reports of low saving rates have repeatedly been revised upward. That is, there is a historical tendency to initially undermeasure the personal saving rate. One reason may be that income is harder to count, and thus easier to underestimate, than is spending, but over time we solve the underestimation problem. So the current situation, in which we are likely undercounting income and saving, is similar to past episodes.

If the personal saving rate is typically understated when first reported, a low level of personal saving may not be very useful for forecasting. Nevertheless, this article will present evidence that *changes* in the saving rate can be used for forecasting movements in income.[1]

First, we need to go back in time to recover the pattern of past reports of income, consumption, and saving, before and after revision. In doing so, we must remember that many initial economic reports rely on data based on surveys that are incomplete and that may contain errors, and the surveys may only imperfectly capture the economic activity they are supposed to target. Over time, more complete data become available as does additional information that helps place each survey in context, making possible a more accurate view of economic activity. As a consequence, economic reports may be revised and may become more accurate. As we shall see, the data on

[1] Here, and elsewhere in the article, income means real income, that is, income adjusted for inflation.

FIGURE 1

(Reasonably) Stable Personal Saving Rate, 1946 to 1992

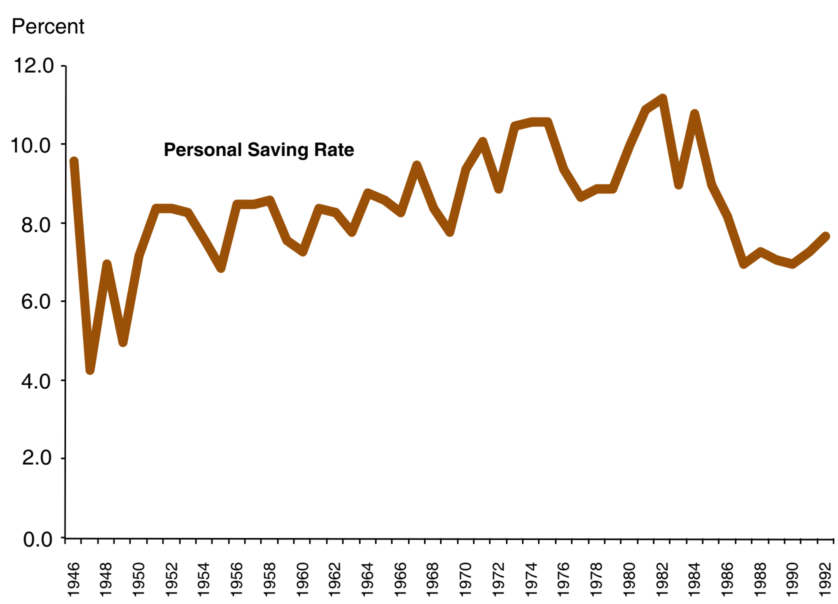

personal saving are particularly vulnerable to changing relationships between initial data and economic activity as a whole. Yet recognizing that the data on saving are imperfect does not exempt us from doing our best to see what information we can extract from the imperfect series.

The inaccuracy of initial economic reports matters because forecasters, and the decision-makers who rely on their forecasts, do not have the luxury of waiting for more accurate measures. They must use the reports available to make their decisions. Economists cannot avoid being concerned about saving and consumption because consumption constitutes a large proportion of demand for output: Personal consumption expenditures have averaged about two-thirds of gross domestic product (GDP) over the last quarter century. To use these data as well as we can, we turn to economic theory, on the one hand, and empirical analysis on the other.

CONSUMPTION THEORY AND FORECASTING

The modern theory of consumption dates back to the 1950s and the work of Milton Friedman. Friedman showed, in the work for which he received the Nobel Prize in 1976, that when our income falls temporarily, we — consumers — are unlikely to reduce our consumption as much as income falls. This argument is called the permanent income hypothesis.

The fundamental argument is that we generally prefer not to consume a lot one year and a little the next; we prefer more equal consumption over the two years. Economists say that consumers prefer a smooth path of consumption rather than one that bounces up and down. In particular, suppose we know that in one year we will have a lot of income and in the following year much less. The preference for smooth consumption means that we will consume about the same each year. So we will save much of our income the first year in order to spend it in the next.[2]

John Campbell, building on sophisticated theories of Friedman's permanent income hypothesis developed by Robert Hall and Marjorie Flavin, has argued that if personal saving fell for consumers as a whole, this would likely forecast an expected increase in income.

The underlying logic can be seen as follows. Suppose consumers raise their consumption while their income remains the same, so their saving falls. Saving can return to normal in one of two ways. Consumers could be intending to reduce their consumption in the future. But that would involve an uneven, rather than a smooth, path for consumption. More likely, consumers have raised their consumption because they expect their incomes to rise in the future.

The Greek fable of the Ant and the Grasshopper, in which a grasshopper who sings all summer starves in the winter, while an ant who saves during the summer is well provided for, serves as a reminder of the possibility that not all households may do a good job of forecasting. Thinking about household behavior raises an important empirical question: How much forethought does the average consumer have? If a large number of households are like the grasshopper, an alternative view of the historically low personal saving rate that the U.S. (and other countries) currently suffers from is that we must inevitably experience a decline in consumption and a recession (Figure 2).

The permanent income theory says that if consumption will be kept the same when income goes up, consumers should expect their future incomes to rise and all of the rise in income to go into saving.

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." In this view, this "downside risk" to spending could trigger a slowdown in economic growth and possibly a recession.

Thus, there are two conflicting notions about how a low saving rate can return to normal. One, the "grasshopper theory," says that consumers will simply consume less and save more. The other, the permanent income hypothesis, says that income will rise while consumption remains stable.

The permanent income theory actually says something more. It says that if consumption will be kept the same when income goes up, consumers should expect their future incomes to rise and all of the rise in income to go into saving. So forecasts of income and saving should tend to mirror one another. To achieve this goal, the forecasting equation for saving must be

[2] A fuller discussion of consumption and the permanent income hypothesis can be found in Satyajit Chatterjee's forthcoming article.

FIGURE 2

Twenty Years of Profligacy? Measured U.S. Personal Saving Rate* As Reported in 2008 Q1

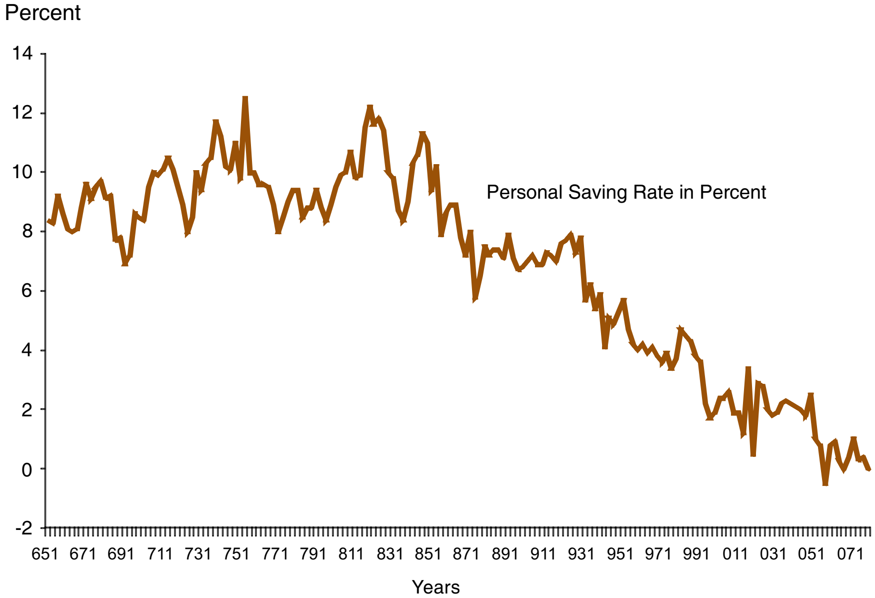

* Seasonally adjusted.

quite similar to the forecasting equation for income. If this tight relationship between the forecasts holds, the two forecasts together should perform better than either separately.[3]

Campbell found evidence for the idea that a low saving rate does imply future growth in income but not for the stronger claim that income growth is exactly related to the size of the drop in the saving rate. However, such exact tests of hypotheses typically fail because hypotheses are necessarily overly simple in their formulation. A statistical rejection can occur because the simple hypothesis is only approximately true, and the data are sufficiently precise to reject the approximation. In fact, Campbell was testing the notion that all consumers are ants, and none are grasshoppers. And that was rejected.

Peter Ireland has pointed out that if Campbell's theory is true, the personal saving rate should be useful in forecasting income, in particular, labor income. He argued that forecasting ability was a good test of an economic theory; indeed, it shows that the economic theory can be useful in a very practical way. That is, the theory says that saving should help us improve our forecasts of income. Because this is true, according to Ireland's argument, the economic significance of the hypothesis is validated, even though statistics may have rejected its narrow implications. For example, when Galileo tested whether two objects of different weight fell at the same speed, he ignored the effect of air resistance. His test, in fact, rejected the hypothesis that the two fell at exactly the same speed. But his test did show that the prediction that they would fall at the same speed was much more accurate than the prediction that they would fall at a speed proportional to their weight. Similarly, although Campbell's estimation showed that forecasted saving did not move with income exactly as the permanent income hypothesis predicts, Ireland's results, as we shall see, showed that when saving and income were assumed to follow the permanent income hypothesis, the forecast was better than if that assumption had not been made. To put it another way, Ireland showed that most consumers were ants rather than grasshoppers, and so, on balance, for practical purposes, there are enough ants that the grasshoppers don't matter.

Before we get to Ireland's evidence, however, we shall first discuss the measurement of saving.

SAVING: INITIAL MEASUREMENT AND REVISION

In a previous *Business Review* article, I argued that the personal saving rate may be mismeasured. The main evidence is that if U.S. personal saving is unusually low, U.S. wealth should be falling. However, the opposite has been true. It is useful to be more precise about how the personal saving rate is measured in the U.S.

The personal saving rate is personal saving as a percentage of disposable (after-tax) personal income. Personal saving, in turn, is disposable personal income minus personal outlays. Disposable personal income includes some easily measured items, such as social insurance contributions

[3] The relationship between the two forecasts is achieved technically through what are called cross-equation restrictions. In this example, a low saving rate forecasts a higher rate of income growth. The low saving rate also forecasts an increase in the saving rate of about the same amount. So the coefficient on the low saving rate should be approximately the same for the two. For details, see the Appendix, and my working paper with Tom Stark.

and benefits. Other parts of labor income, such as other (that is, non-Social Security) benefits and transfers, are subject to measurement and conceptual problems. (For example, is a pension considered income when it is earned or when it is paid to a worker?) Wages and proprietors' income are subject to underreporting in government records as a result of tax evasion. 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.

Under the current method of measuring income, we may not be capturing all of the sources of household income, and this may result in the appearance of low saving. In that case, if, in the future, we figure out a better way to measure income, we will revise our current estimates of saving upward.[4]

How are data revised? Data on a given quarter's economic activity are first published in an advance estimate, late in the first month of the next 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. Initial estimates thus undergo three summer revisions. Thereafter, the estimates are changed only in benchmark revisions, which now occur every four years. Benchmark revisions provide an opportunity for the Bureau of Economic Analysis (BEA) to make discretionary choices in defining the items it considers to be part of personal income. For example, government pension income is now considered income when it is earned, rather than when it is paid out.[5] In addition, more complete data from economic censuses are included at this time. Most of the revisions to the saving rate that turn initially low rates into higher ones occur during benchmark revisions.

Real-Time Data Collected by the Philadelphia Fed. Researchers Dean Croushore and Tom Stark pioneered the collection of data sets in vintages that capture the data as they were available on a particular date. These data can be used to show how revisions change our view of economic processes. (The real-time data used here as well as a number of other statistical series can be downloaded at the following address: http://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data/.)

These data show that adjustments to the measurement of personal saving have occurred repeatedly in the past. For example, the average saving rate from 1980 to 1985 was initially reported as 6.5 percent and is now reported as 10.4 percent, the highest saving rate in our current data series. The advance estimates of the personal saving rate from the third quarter of 1965 to the second quarter of 1999 averaged 5.3 percent. But after revisions, as reported in September 2007, the personal saving rate over the same period averaged 8.1 percent. Figure 3 shows the advance

FIGURE 3

Before and After: Measured Personal Saving Rate as First Published and as of 2007 Q4

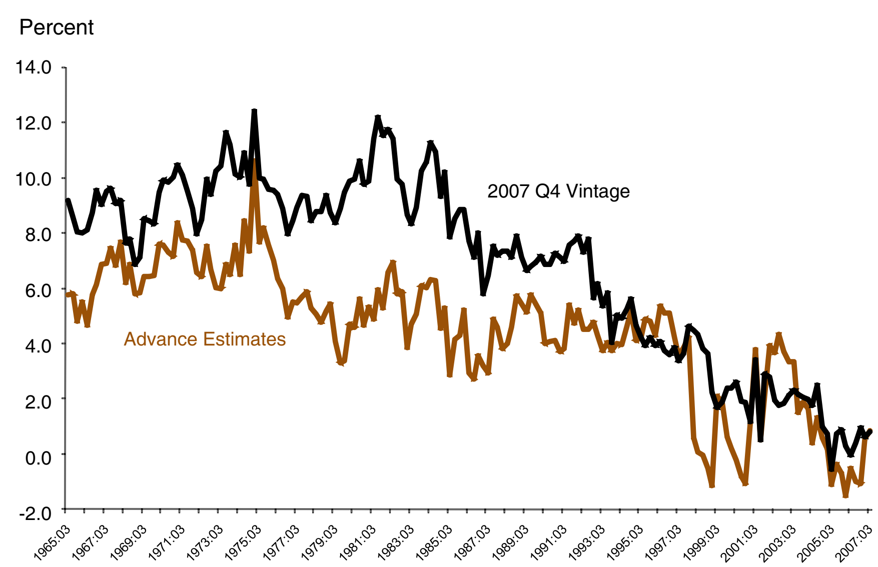

[4] In particular, it's possible that we will eventually adjust our treatment of capital gains, which have become an important contributor to the increase in household wealth.

[5] When government pensions were mainly federal pensions, they were treated this way because the federal government did not set aside income at the time the pension obligations were incurred. However, now most government pensions are state and local government pensions, and these governments generally set aside pension funds when their employees earn the pensions.

estimates of the personal saving rate as they were reported when first published, and the latest vintage data, as we would have seen them last year. For example, in the fourth quarter of 1985, when the advance report for the personal saving rate in the third quarter of 1985 became available, it was reported as 2.9 percent. (See the table in *Understanding and Using Real-Time Data*.) But if we go to the BEA's website today, we will find that we now believe that the personal saving rate in the third quarter of 1985 was 7.9 percent.

What is the problem in measuring saving? It turns out that complete data on income are hard to measure. As the economy evolves, new types of income come about. Initially, the new income may not be reported or may not be considered income. Over time, as new sources of data become available and as old data come to be viewed in a new light, more income is reported.

These changes in income are usually recorded in the benchmark revisions. In our working paper, Tom Stark and I show that almost all of the upward revisions to the personal saving rate occur in benchmark revisions.

Which data should we use for our tests of the permanent income hypothesis? Peter Ireland and Tom Stark and I simulated forecast exercises. Ireland's test focused on how consumers actually behave, to test whether the underlying consumer behavior was primarily driven by households that conform to the permanent income hypothesis. To do this test, we want to use the data that best reflect the underlying behavior, that is, the most accurate available data. And those are the latest revised data, in the most recent vintage.

Forecasting well is not necessarily the same as understanding consumer behavior. For understanding consumer behavior, how poor our latest statistics are is irrelevant — what we care about is the underlying behavior revealed by the best statistics, which may be available only in historical data that have been revised. Ireland used the right data from the perspective of understanding consumer behavior — the best and latest available statistics — but those data are not the best guide for understanding how to forecast with the data the forecaster actually has available.

Forecasting well is not necessarily the same as understanding consumer behavior.

If we want to test how useful the saving rate is to an economic forecaster, we should use real-time data, which will repeatedly put us into the situation of the forecaster: using data that have not yet been revised.

FORECASTING WITH MISMEASURED PERSONAL SAVING

One way economists analyze statistical relationships in economic data is to perform an in-sample data analysis called a regression analysis. In our case, we want to examine whether in periods when saving is low, income rises faster than usual in later periods. This statistical relationship can then be used to forecast the behavior of the consumer.

Peter Ireland's insight was to argue that even if not all consumers behave according to the permanent income hypothesis, if most of them do so, it may be better to assume than not assume the hypothesis that saving and income will rise together and that assumption will produce better forecasts.

The method Ireland used for his test is called *recursive* out-of-sample testing. This method basically asks over and over (recursively) whether the relationship in past data successfully forecasts the next piece of data. This analysis is out-of-sample because the next piece of data is never in the sample.

To use a specific example, we take a base period, say, from the first quarter of 1959 to the fourth quarter of 1970, and do a regression analysis. We then use this regression analysis to forecast the next quarter's income — the first quarter of 1971. We then compare this to the actual income for the first quarter of 1971 and measure the error in this forecast. We then add the first quarter of 1971 to our data, lengthening our data one period, and undertake a new regression based on data from the first quarter of 1959 to the first quarter of 1971. We then forecast the new next period: we forecast income in the second quarter of 1971 and again measure the error in this forecast.

Continuing to the present, we can accumulate a long series of forecasts, the actual data, and the forecast errors.[6] We square the errors, sum them up, and divide by the number of forecasts to obtain the mean square error of the forecasts. We then take the square root to obtain the root mean square error, a number conceptually similar to the standard deviation. The smaller the root mean square error, the more accurate the average forecast. When we make forecasts of income using past income and saving, we will compare the root mean square error with the root mean square error when only past income is used in the

[6] The forecast error is the difference between the actual value and the predicted (forecast) value of a time series.

Understanding and Using Real-Time Data

Ordinarily, data used in economic analysis are what real-time data users call the latest available vintage. It reflects the data that were published by the statistical agency (in the case of the U.S. personal saving rate, the Bureau of Economic Analysis, or BEA) at the time the economic analysis was performed. As we have stressed, these data may look very different from those that were available to a forecaster at some earlier time.

The table on page 16 contains selected portions of a real-time data matrix. Each column in the full matrix represents a "vintage." Each vintage contains the data from the first quarter of 1947 to the quarter before the vintage date, as it was published at the vintage date.

A forecaster in the fourth quarter of 1985 would have had available the data in column 2 of the table presented here and would have thought that the personal saving rate in the third quarter of 1985 was 2.9 percent. A forecast of real income in the fourth quarter of 1985 would have been based on this estimate.

In the next quarter, the BEA published a benchmark revision, and the personal saving rate for the third quarter of 1985 then appeared to be 3.7 percent. This is shown in column 3. Because this is a benchmark revision, the entire history of the personal saving rate has been revised. Note that even the data from 1947 have been revised. This would presumably have caused a forecaster using the personal saving rate in forecasting to redo the regression analysis on which the forecast was based. Again, in the third quarter of 1986, a summer revision changed the data for the past three years, and the estimate of the personal saving rate for the third quarter of 1985 was reported to be higher yet, 4.2 percent (column 5).

The data that Peter Ireland used in his paper were those published by the BEA in the fourth quarter of 1994 (column 6). In Ireland's work, the estimate of personal saving for the third quarter of 1985 was 5.4 percent. For all of his forecasts, Ireland would have used column 6 data. In our real-time forecast analyses, presented here, we use a different column of data for each forecast. Thus, we would assume that the forecaster in the fourth quarter of 1985 would use the data in column 1 to estimate the forecast equation and to make the forecast. In the first quarter of 1986, the forecaster would use the data in column 2. By contrast, Ireland's forecaster in the fourth quarter of 1985 uses the data in column 6, up to the row that says third quarter of 1985, to make a forecast.

Remarkably, the data change further in the latest vintage used in our study, the third quarter of 2005, which shows a personal saving rate of 7.9 percent in the third quarter of 1985 (that's the number still being reported as of the first quarter of 2008).

An interesting contrast is to look at changes in the personal saving rate, which changes much less across vintages compared with the level. In both the real-time vintage of the fourth quarter of 1985 (column 2), and in the "latest" vintage of the third quarter of 2005 (column 7), we observe that the personal saving rate in the third quarter of 1985 is low relative to its neighbors in the quarter before and the quarter after. In our working paper, Tom Stark and I present additional evidence that changes in the personal saving rate are more stable over time than the level, which is an important reason why changes in the saving rate have better predictive power than the level of the saving rate.

forecast. If saving does help to forecast income, it will lower the root mean square error.

Peter Ireland used this "recursive regression" method to forecast income using the personal saving rate.[7] Using data from 1959 to 1994, he showed that the personal saving rate was a good forecaster of income from 1970 to 1994.

To create a benchmark for his forecasts, he began by using past values of income growth to forecast future income growth. There is good evidence that — in part because economies tend to go through booms and busts — when income growth is high, it tends to remain high, and when income growth is low, it tends to remain low, a pattern called persistence. Ireland made recursive out-of-sample forecasts of income growth using past values

[7] When the regressions are formed by continually enlarging the data set, so that, as in the example, we always begin from 1959, the regressions are called recursive. An alternative technique is "rolling regressions," where, as we add more recent data, we drop off the oldest data, so that the period under consideration is always the same length.

TABLE

Example of Real-Time Data Personal Saving Rates in Six Vintages, Selected Observations

1	2	3	4	5	6	7
	Vintage 85:Q4	Vintage 86:Q1	Vintage 86:Q2	Vintage 86:Q3	Vintage 94:Q4	Vintage 05:Q3
Date						
1947:Q1	5.0%	4.9%	4.9%	4.9%	4.8%	6.1%
1947:Q2	1.4%	1.3%	1.3%	1.3%	1.2%	2.6%
1976:Q1	7.7%	8.2%	8.2%	8.2%	7.9%	9.6%
1976:Q2	7.3%	8.0%	8.0%	8.0%	7.7%	9.6%
1976:Q3	6.7%	7.6%	7.6%	7.6%	7.3%	9.5%
1976:Q4	5.9%	6.9%	6.9%	6.9%	6.6%	8.9%
1984:Q1	6.1%	7.0%	7.0%	6.9%	8.1%	10.3%
1984:Q2	5.7%	6.1%	6.1%	6.0%	7.8%	10.6%
1984:Q3	6.3%	6.7%	6.7%	6.4%	8.4%	11.3%
1984:Q4	6.2%	6.0%	6.0%	6.0%	7.9%	11.0%
1985:Q1	4.5%	4.8%	4.8%	5.2%	6.7%	9.4%
1985:Q2	5.1%	5.9%	5.9%	6.5%	7.8%	10.2%
1985:Q3	2.9%	3.7%	3.7%	4.2%	5.4%	7.9%
1985:Q4	#N/A	4.1%	4.0%	4.4%	6.0%	8.6%
1986:Q1	#N/A	#N/A	4.3%	5.0%	6.5%	8.9%
1986:Q2	#N/A	#N/A	#N/A	5.2%	7.2%	8.9%
		↑ Benchmark Revision		↑ Summer Revision	↑ Ireland Data	↑ Nakamura-Stark Latest Available Vintage

of income growth and then measured the forecast error. He repeated this over the period from 1970 to 1994 and calculated the root mean square error.

He then made similar forecasts of future income growth using past values of income growth and adding past values of saving. He found that the root mean square error was lower than when only past values of income were used. Moreover, he found that the forecast error was even lower when he accounted for the restrictions imposed by the permanent income hypothesis: that predicted savings and income have parallel movements. He took this to be good evidence that the permanent income hypothesis is true.

However, Ireland used the data as they were available in 1994. This is not really a true test of personal saving's usefulness in forecasting because we know that the personal saving rate as it was available in 1994 differed substantially from what it looked like in, say, 1980. So Ireland, making his forecasts in 1994, used an estimate of the personal saving rate for the third quarter of 1985, for example, that was 5.4 percent, while the forecaster in the fourth quarter of 1985 would have thought it was 2.9 percent (and, as we now know, it was later revised to 7.9 percent).

FORECASTING WITH REAL-TIME DATA

Using real-time data from the Philadelphia Fed's data set, we can make real-time forecasts that use the data as they were available to an economist on a series of dates. (To see how these data are organized, see *Understanding and Using Real-Time Data*. Further information can be found in the 2000 article by Croushore and Stark.) Real-time data enable us to ask: Given that personal saving has historically been dramatically mismeasured, would it be a useful forecasting tool?

Forecasting Income with Saving, with Latest Available Data, and in Real Time. With data that have been revised over many years, the relationship between the level of saving and future income growth is just as the permanent income hypothesis shows, as Peter Ireland also showed.

However, if we try to do the same exercise in real time, the level of the saving rate is not predictive. I will show that, in particular, from 1981 to 2005, the level of the saving rate does not improve forecasts of income growth. All is not lost, however, because I will show that changes in the saving rate can be used in real time to forecast income growth.

Forecasting in Real Time. Let's look at the forecasts using real-time data, shown in the first row of the table on page 18. If we look at the period before 1982 (the first quarter of 1971 to the fourth quarter of 1981), before the saving rate started trending downward, even in real time there is value to these forecasts, although the improvement shrinks to 3.7 percent. There is even a small improvement from imposing the restrictions of the permanent income hypothesis.

But when we look at the data after 1981, the level of the saving rate is much less helpful in forecasting. Looking at row 2, from the first quarter of 1982 to the second quarter of 2005, we see that when the level of the saving rate has been falling, it has been much less useful in forecasting. In real time during this period, the level of the saving rate worsens forecasts, with or without the restrictions. As we see in row 2, the forecasts are 4 percent worse using the level of the saving rate and 1 percent worse adding the restrictions.

An Alternative: Forecasting with the Change in the Saving Rate. Thus far, I have focused on the level of the saving rate as a measure of future income expectations because the underlying theory and the data suggest that the level of the saving rate should generally be stable. Therefore, when the saving rate is below average, we expect it to rise toward the average. A below-average level of the saving rate, according to the permanent income hypothesis, implies that income is expected to rise, causing saving to rise. But, as we have seen, the most recent level is typically too low and likely to be revised higher. Thus, the level might be misleading. Perhaps we should try the change in the saving rate. Even if the level is low because of mismeasurement, a downward change might be telling us that income is expected to increase.

It is true that in the absence of substantial measurement error, the change in the saving rate is unlikely to be as informative as the level of the saving rate. If we look at the latest revised data, in the heavily revised period from 1971 to 1981, we see that the level of the saving rate reduces the root mean square error 12 percent. When

Using real-time data from the Philadelphia Fed's data set, we can make real-time forecasts that use the data as they were available to an economist on a series of dates.

TABLE

Forecasting Real Disposable Income Growth with Real-Time Data:
Ratios of Forecast Errors, Forecasts with Saving Relative to Forecasts with Only Past Income Growth*

	(1) Level of Saving	(2) Permanent Income Hypothesis Restrictions on Level of Saving	(3) Change in Saving	(4) Permanent Income Hypothesis Restrictions on Change in Saving
1971:Q1 – 1981:Q4				
1. Real time	0.963	0.954	0.950	0.944
1982:Q1 – 2005:Q2				
2. Real time	1.040	1.010	0.943	0.935

* Lags chosen using the Akaike information criterion.

we add the restrictions of the permanent income hypothesis, we reduce the root mean square error 16 percent. With the latest revised data, the change in the saving rate does not do as well in this period, reducing the root mean square error 10 to 12 percent, depending on whether we impose the restrictions of the permanent income hypothesis.

Thus, with good revised data, the change in the saving rate is not as informative about future changes in income as is the level of the saving rate. The theory points us to the right form for the data.

But as noted before, this does not tell us about the situation a forecaster faces. If we look at that same period but make forecasts using real-time data, we see that the level of the saving rate reduces the root mean square error only 3.7 percent, and the permanent income hypothesis restrictions add only a small improvement, reducing the root mean square error 4.6 percent. We make better forecasts with the change in the saving rate, which produces a 5 percent improvement without the restrictions of the permanent income hypothesis and 5.6 percent with them.

If we look at the more recent period, from 1982 to 2005, the level of the saving rate performs quite poorly in forecasting. In real time in this period, the level of the saving rate worsens forecasts with or without the permanent income hypothesis restrictions, as we have seen. By contrast, the change in the saving rate performs well, reducing the root mean square error 5.7 percent without the permanent income hypothesis restrictions and 6.5 percent with it.

Thus, using changes in the saving rate in real time, a forecaster could have made a better forecast of future income than using only past data on income. This is true whether or not the permanent income hypothesis restrictions are imposed. By contrast, the level of the saving rate, despite attractive theoretical properties and despite the fact that the level of the saving rate does well with the latest revised data overall, would not have been a good choice in a forecasting equation over the past 20 years.

Why might the change in the saving rate be better in real time than the level of the saving rate? It turns out that the change in the saving rate is subject to smaller revisions than the level of the saving rate. Technically, this is because revisions tend to have a cumulative impact on the levels. Consequently, the changes are more reliable than are the levels.

CONCLUSION

I have made three points in this article. First, I argued that when the saving rate falls, it is more likely to be evidence that households expect faster real income growth in the future, rather than evidence that they are spending too much and will have to cut back on consumption.

Second, I showed that the personal saving rate has typically been substantially revised and usually upward. The Philadelphia Fed's real-time data set gives us the data we need to show this. Since a low personal saving rate can occur because of mismeasurement and may well be revised upward, in practice, the level of the personal saving rate does not help us forecast real income growth.

Finally, I showed that, guided by this insight, forecasters can use the change in the saving rate rather than the level as a forecasting tool. Although this technique does not work as well as having better data would, it does enable economists to improve their forecasts.

So I have shown that real-time data can be quite useful for improving forecasting when revisions are large. By using real-time data, economists can sometimes figure out how current data can be valuably employed, even when poorly measured.

TECHNICAL APPENDIX

Permanent Income Hypothesis Restrictions

Campbell's version of the permanent income hypothesis that we are testing is a two-equation system that predicts changes in income and the level of the saving rate. The system relates these to past values of income and the level of the saving rate.

Formally, the system is

$$\Delta Y_{lt} = a(L)\Delta Y_{lt-1} + b(L)S_{t-1} + u_{1t}$$
$$S_t = c(L)\Delta Y_{lt-1} + d(L)S_{t-1} + u_{2t}$$

where Y_{lt} 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_t are forecast error terms.

Whenever the expected permanent increase in real labor income occurs, the saving rate is expected to rise at the same time. The permanent income hypothesis says that these two expected increases are closely related; econometrically, this relationship is called a cross-equation restriction because it relates coefficients across the two equations. The intuition behind these cross-equation restrictions is that a current decrease in saving must imply a future predictable permanent increase in real labor income and a future predictable saving increase. The 2p restrictions on the coefficients of the lag operators are

$$c_i = a_i, i = 1, \ldots, p$$
$$d_1 = b_1 + (1+r)$$
$$d_i = b_i, i \geq 2,$$

where r represents a constant real interest rate. See Peter Ireland's article for a more detailed description.

REFERENCES

Campbell, John Y. "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Personal Income Hypothesis," *Econometrica* 55 (November 1987), pp. 1249-73.

Chatterjee, Satyajit. "The Peopling of Macroeconomics: Microeconomics of Aggregate Consumer Expenditures," Federal Reserve Bank of Philadelphia *Business Review* (forthcoming).

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 Funny Thing Happened on the Way to the Data Bank: A Real-Time Data Set for Macroeconomists," Federal Reserve Bank of Philadelphia *Business Review* (September/October 2000), pp. 15-27.

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), pp. 605-17.

Croushore, Dean, and Tom Stark. "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics* 105 (November 2001), pp. 111-30.

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

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

Ireland, Peter N. "Using the Permanent Income Hypothesis for Forecasting," Federal Reserve Bank of Richmond *Economic Quarterly* 81 (Winter 1995), pp. 49-63.

Kilian, Lutz, and Atsushi Inoue, "In-Sample or Out-of-Sample Tests of Predictability, Which One Should We Use?" *Econometric Reviews* 23 (November 2004), pp. 371-402.

Nakamura, Leonard. "Investing in Intangibles: Is a Trillion Dollars Missing from GDP?," Federal Reserve Bank of Philadelphia *Business Review* (Fourth Quarter 2001).

Nakamura, Leonard, and Tom Stark, "Mismeasured Personal Saving and the Permanent Income Hypothesis," Federal Reserve Bank of Philadelphia Working Paper 07-8 (February 2007); available at: http://www.philadelphiafed.org/research-and-data/publications/working-papers/2007/wp07-8.pdf

Stark, Tom, and Dean Croushore. "Forecasting with a Real-Time Data Set for Macroeconomists," *Journal of Macroeconomics* 24 (December 2002), pp. 507-31.