Consumption-based macroeconomic forecasting

Jeffrey R. Campbell and Spencer D. Krane

Introduction and summary

Macroeconomics explains the behavior of broad measures of economic activity, such as the total amount of income produced in the economy or the overall average of the prices households pay for the goods and services that they consume. Such aggregates reflect the confluence of many different decisions in product and financial markets: Households decide how much to work and what to spend; businesses commit to capital investment projects; foreigners decide on their demand for U.S. products; and governments determine taxes and spending.

To discern how particular events influence macroeconomic performance, researchers and policymakers sometimes turn to large-scale econometric models. These are collections of statistically estimated equations, which describe a wide range of economic decisions. The results from these equations are then aggregated into macroeconomic outcomes. Large-scale models can address a wide range of relevant questions, but their size comes with a cost. To make estimation and aggregation practical, they require numerous restrictions. Sims (1980) argued that these restrictions often lack theoretical justification, which can cast doubt on some of their predictions. Furthermore, the complexity of large models sometimes makes it difficult to determine the relevant factors underpinning their results.

An alternative approach uses the behavior of broad macroeconomic aggregates to identify a few fundamental sources of change in aggregate production and prices. We follow this methodology and build a small-scale econometric model designed to separate the influence of permanent and transitory factors on the level of economic activity. This distinction can usefully inform policymakers' decisions. For example, in Taylor's (1993) formulation of monetary policy rules, the interest rate responds to a transitory deviation in output from its long-run potential level but not to simultaneous movements in both actual and permanent output.

Because the model is small and makes few assumptions, it cannot identify a large number of independent factors that may affect economic outcomes. However, its simplicity is also a virtue—the model's results have simple interpretations and are robust to a number of specification issues that can plague large-scale systems.

Even small econometric models like ours require assumptions to identify the sources of economic fluctuations. For this purpose, we employ Friedman's (1957) permanent income theory of consumption. The central prediction of this theory is that forward-looking households consume a constant fraction of their permanent income—the sum of their financial wealth and the financial value of their current and *future* labor earnings. This leads us to use the identifying assumption that all permanent changes in nondurable consumption reflect changes in households' permanent income.

The model links nondurable consumption with other economic expenditures using the observation of *balanced growth*: households' expenditures on very broad categories of goods and services do not drift apart over time. With balanced growth, permanent changes in the consumption of nondurable goods and services eventually lead to equally sized changes in other spending. The model we estimate imposes balanced growth across consumption of nondurable goods and services, durable goods expenditures, and total private savings.[1]

We estimate our model using data on the U.S. economy from 1983:Q1 through 2005:Q2. This sample deliberately omits the period of high and variable inflation beginning in the early 1970s, so the model's

Jeffrey R. Campbell is a senior economist and Spencer D. Krane is a vice president and economic advisor in the Economic Research Department at the Federal Reserve Bank of Chicago. *The authors thank Craig Furfine and Marcelo Veracierto for their helpful comments and suggestions and Nishat Hasan for research assistance.*

dynamics reflect the monetary policy from the period of low and stable inflation that followed Paul Volcker's chairmanship of the Federal Reserve. Furthermore, since the early 1980s the volatility of many macroeconomic aggregates has been much lower than it was in the immediate post-Korean War era. Our model will reflect the structure of the economy during only this less volatile period. Of course, using such a short sample limits the number of parameters that we can estimate. This also leads us to highly value parsimony.

The remainder of this article proceeds as follows. We next review the economic theory that guides our forecasting model's specification. Then we examine the implications of this theory for our empirical model and describe the data. After this, we present the estimated model and its forecasts. Finally, we discuss the use of our model to measure and forecast the gap between actual output and its long-run potential level.

Theoretical foundations

Theoretically oriented macroeconomics guides the specification of our forecasting model, so we begin our analysis with a review of two of its important results. First, consumption of nondurable goods and services should be helpful for forecasting other macroeconomic quantities. Second, the shares of households' expenditures on broad classes of items do not exhibit permanent changes.

We derive both of these results from the theory of a forward-looking household's optimal consumption and savings decisions, developed originally by Friedman (1957). The household values two goods, consumption of nondurable goods and services, C_t, and the service flow from a stock of durable goods, S_t. For simplicity, we assume the service flow from durable goods is equal to the stock.[2]

The household's resources are its labor income and its initial wealth. Labor income in period t is Y_t, and this is potentially random. The household places all of its financial wealth in a single asset, risk-free bonds. The face value of the bonds at the beginning of period t is B_t, and the household can purchase bonds that come due in period $t + 1$ at the price $1/R_t$. We assume that each period is a calendar quarter, so the implied annual interest rate on these bonds in percentage points is $400 \times (R_t - 1)$. Both bonds and labor income are denominated in units of the nondurable good. The household can trade these resources for durable goods at the relative price P_t. Let X_t and δ represent the household's purchases of durable goods and the constant rate at which they depreciate so that $S_t = (1 - \delta)\, S_{t-1} + X_t$.

Given the household's initial financial wealth (B_0) and the value of its used durable goods stock ($P_0(1 - \delta)\, S_{-1}$), the household allocates its resources across the consumption of nondurable goods and services, durable goods purchases, and bonds in order to maximize the expected utility function,

$$E_0\left[\sum_{t=0}^{\infty}\beta^t \ln\left(C_t^{\theta} S_t^{1-\theta}\right)\right],$$

subject to the sequence of budget constraints,

$$1)\quad C_t + P_t X_t + B_{t+1}/R_t = Y_t + B_t.$$

Here, $E_0[Z_t]$ denotes the mathematical expectation of the random quantity Z_t, calculated using information available at time 0.[3] Increasing β increases the household's value of future consumption relative to current consumption. In this sense, β measures the household's patience. The exponent θ lies between 0 and 1: This ensures that utility increases and marginal utility falls with increasing consumption of either C_t or S_t. The larger θ is, the more the household values nondurable goods relative to durable goods.

The role of expected future income in current consumption decisions can be most easily appreciated by replacing the sequence of budget constraints (equation 1) with a single unified budget constraint,

$$2)\quad \sum_{t=0}^{\infty}\hat{R}_t^{-1}E_0\left[C_t + P_t X_t\right] = B_0 + \sum_{t=0}^{\infty}\hat{R}_t^{-1}E_0\left[Y_t\right].$$

Here, $\hat{R}_t$ is the interest rate on a t-period risk-free bond sold in period 0. The left-hand side equals the present value of the household's expenditures on nondurable and durable consumption goods, while the right-hand side equals the present value of its assets—its financial wealth plus the value of its current and future labor income.[4] The feature of equation 2 that is important for our purposes is that an increase in $E_0[Y_t]$ increases the resources available for consumption at *all* dates, including the present.

Consumption as a predictor of future income

To demonstrate the forecasting power of nondurable consumption for income, we begin by characterizing the household's savings decision. If C_t and B_{t+1} are optimal, then the household cannot make itself better off by reducing C_t slightly, using the foregone consumption to purchase bonds and consuming the principal and interest in period $t + 1$. The utility cost of slightly decreasing C_t is θ/C_t, and the future utility benefit from temporarily increasing savings is $R_t\theta/C_{t+1}$. This benefit is potentially random, because the household's actual choice of C_{t+1} could be affected by random changes in labor income or interest rates in period $t+1$. Hence, we wish to compare the cost of foregone consumption with the *expected* future benefit, $E_t[R_t\theta/C_{t+1}]$. Discounting this back to period t and equating it to the current-period utility cost yields

$$3)\quad \frac{1}{C_t} = \beta R_t E_t\left[\frac{1}{C_{t+1}}\right].$$

Because the consumer's choices of C_{t+j} and B_{t+j+1} must also be optimal, given the information available at time $t+j$, a version of equation 3 must hold in all future periods. That is,

$$4)\quad E_t\left[\frac{1}{C_{t+j}}\right] = \beta E_t\left[R_{t+j}\frac{1}{C_{t+j+1}}\right].$$

Equation 4 arises from taking expectations of both sides of equation 3 after first increasing the dates of all variables by j periods.

Equation 3 implies that changes in the household's expectations of future income directly influence C_t. To see this, hold the interest rate at some constant value and suppose that $E_0[Y_t]$ increases so that the right-hand side of equation 2 rises by 1 percent. If the household spends all of this additional income on nondurable goods, then the only solution to the optimality conditions, equation 3 and equation 4, is to increase C_t and all of its future values by 1 percent. Allowing the household to use some of this extra income to purchase durable goods changes the magnitude of the consumption response, but not the result that the current and future percentage responses are the same. In this sense, C_t is a forward-looking variable that should be informative about the household's expectations.

A second implication of equation 3 is that no macroeconomic variable can improve a forecast of C_{t+1} that already uses C_t. This can be seen by multiplying both sides of equation 3 by C_t, setting $R_t = \beta^{-1}$, and rearranging terms to get

$$5)\quad E_t\left[\frac{C_{t+1}-C_t}{C_{t+1}}\right] = 0.$$

Equation 5 embodies Hall's (1978) result that no information available at time t is useful for forecasting the growth rate of nondurable consumption. If we relax the strong assumption that $R_t = \beta^{-1}$, then the appropriately modified version of equation 5 implies that the interest rate is the only variable with information about the growth rate of consumption. Together, these results characterize the role of C_t in macroeconomic forecasting: Nondurable consumption is informative about future income, but only the interest rate can help predict its growth.

Durable goods

To tie the evolution of nondurable consumption with other macroeconomic aggregates, we begin by characterizing the household's utility-maximizing durable goods purchases. Again, consider a small change to the household's optimal expenditures. Suppose that the household sells P_t units of nondurable consumption for one unit of durable goods consumption, holds that durable good until the next period, and then sells it on a used durable goods market. Since the household's initial choices were optimal, this small adjustment cannot increase utility. The cost of this adjustment is $P_t\theta/C_t$, the utility value of the foregone nondurable consumption. There are two benefits. First, the consumer enjoys the utility of the service flow from the additional durable goods, $(1-\theta)/S_t$. Second, next quarter the consumer can expect to sell the depreciated durable goods for $(1-\delta)P_{t+1}$ units of nondurable consumption goods which are worth marginal utility of $\beta\theta/C_{t+1}$. Equating the costs with the expected benefits, assuming that P_t grows at the constant rate γ, and rearranging terms yields

$$6)\quad \frac{P_t S_t}{C_t} = \frac{1-\theta}{\theta}\left(\frac{R_t-(1-\delta)\gamma}{R_t}\right)^{-1}.$$

On the right-hand side of equation 6, the term $(R_t - (1-\delta)\gamma)/R_t$ can be interpreted as the upfront cost of renting $1/P_t$ units of durable goods. This equals the period $t+1$ expense of repaying the interest and principal from borrowing the purchase price minus the resale value discounted back to period t. If the interest rate, preferences for durable goods (the parameters β and θ), and technology (δ and γ) remain unchanged, then this rental cost is constant. In this case, the right-hand side of equation 6 is constant, so the ratio of nominal expenditures on nondurable consumption to the value of the durable goods stock also does not change over time. This is one aspect of balanced growth, which we now proceed to examine in more detail.

Balanced growth

We call economic growth *balanced* if it leaves households' expenditure shares on broad classes of items unchanged in the long run. Balanced growth ties the long-run levels of macroeconomic quantities together, thereby aiding forecasting.

From the end of World War II until the early 1980s, households' financial savings, their purchases of durable goods, and their expenditures on nondurable goods and services were balanced. This motivated the builders of early general equilibrium business cycle models, such as King, Plosser, and Rebelo (1988), to assume that as income rises, households will continue to consume the same shares of the various goods and services available in the economy.

We present here the theoretical foundations of balanced growth. Suppose that the household's labor income grows at the constant rate $Y_{t+1}/Y_t = \mu$, the interest rate equals the constant $R = \beta^{-1}\mu$, and the price of durable goods equals P always. We wish to find the utility maximizing choices of C_t, S_t, and B_{t+1} given the household's initial labor income, Y_0, and tangible wealth, $B_0 + P(1-\delta)S_{-1}$.

On any balanced growth path, $C_{t+1}/C_t = S_{t+1}/S_t = B_{t+1}/B_t = \mu$.[5] It turns out that the unique choice of C_0 that is consistent with this sequence satisfying $B_{t+1}/B_t = \mu$ is

7) $C_0 = \kappa(Y_0/(1-\beta) + B_0 + P(1-\delta)S_{-1})$.

Here, κ replaces a complicated expression of β, θ, μ, and δ. The sum multiplying κ is the value at time 0 of the household's permanent income. This is the sum of its tangible wealth and the present value of its current and future labor income, $Y_0/(1-\beta)$.[6]

We are also interested in the growth rates of durable consumption expenditures and household savings. Durable goods purchases directly inherit the growth rate of the durable goods stock, because

$$\frac{X_{t+1}}{X_t} = \frac{S_{t+1}}{S_t}\frac{\left(S_{t+1}-(1-\delta)S_t\right)}{S_{t+1}}\frac{S_t}{\left(S_t-(1-\delta)S_{t-1}\right)} =$$

$$\mu\left(1-(1-\delta)/\mu\right)\frac{1}{1-(1-\delta)/\mu} = \mu.$$

Net private savings equals $Y_t + (R-1)B_t - C_t - PX_t$. All four terms grow at the rate μ, so their sum does as well. Thus, the model predicts that the household's expenditures on nondurable goods, its durable goods purchases, and its savings all grow at a common rate, μ.

Permanent income shocks

We interpret the balanced growth path as a description of the household's long-run choices in the absence of business cycles. However, the solution in the previous subsection can be used to develop intuition about short-run responses to changes in permanent income. To do so, suppose that the household is originally on a balanced growth path so that the solution in equation 7 sets $C_0/C_{-1} = \mu$. We say that the household receives a permanent income shock if Y_0 unexpectedly rises above μY_{-1} but all other growth rates of Y_t remain unchanged. Equation 6 requires both C_0 and S_0 to increase by a common percentage. Typically this requires the household to spend some of its savings. Thereafter, C_t, S_t, and B_t all continue to grow at the rate μ.

Consider the observable implications of this response for consumption expenditures. The behavior of C_t following a permanent income shock is consistent with the analysis in the previous subsection on consumption as a predictor of future income: The shock to Y_0 causes a one-time shift up in the path for C_t and induces no forecastable changes in its subsequent growth rates. Next, consider the shock's impact on X_t. Purchases in period $t = 0$ must cover both depreciation on the existing stock and the increase in the desired stock to S_0; that is, $X_0 = S_0 - (\mu-1)S_{-1} + \delta S_{-1}$. Rearranging terms in this expression allows us to relate the growth rate of S_0 with that of X_0.

$$\frac{X_0 - X_{-1}}{X_{-1}} = \left(\frac{S_0 - S_{-1}}{S_{-1}}\right)/\left(\mu - 1 + \delta\right).$$

For empirically relevant choices of μ and δ, the term $\mu - 1 + \delta$ is considerably less than one. Thus, the given change in S_0 translates into a much larger percentage change in X_0. This change is mostly temporary so that $X_1 = (\mu + \delta)S_1$. Hence, the model leads us to expect expenditures on durable goods to be much more volatile than expenditures on nondurable goods. This volatility reflects transitory aspects to growth that arise from the role of durable goods expenditures as an investment in household capital.

The forecasting model

The remainder of this article describes the construction and use of a vector autoregression (VAR) forecasting model of the U.S. economy. The three key aggregates from the U.S. *National Income and Product Accounts* (NIPA) that it includes are privately demanded gross domestic product (GDP), personal consumption expenditures on nondurable goods and services, and personal consumption expenditures on household durables. The model builds on the theory of household consumption and savings in two ways. First, it uses consumption expenditures on nondurables and services to identify shocks to permanent income. Second, it constrains the evolution of durable consumption expenditures and private savings so that the model's forecasts satisfy balanced growth restrictions in the long run.

The theory identifies the interest rate of a consumption-denominated bond as informative for both consumption growth and durable goods purchases, so the forecasting model also uses information on the real cost of borrowing. The interest rates in U.S. financial markets are typically denominated in dollars, so they must be adjusted by market participants' expectations of inflation to produce a real interest rate directly relevant for households' decisions. Instead of explicitly measuring inflation expectations and using these to directly construct the real interest rate, we include the interest rate on federal funds and the inflation rate in the forecasting equations of all variables. If market participants forecast inflation with some linear combination of current and lagged variables in the model, then we can interpret the estimated equations as including market participants' inflation forecasts. Inflation and interest rates are themselves of substantial macroeconomic interest, so the model includes forecasting equations for them as well.

FIGURE 1

Growth of nondurable consumption

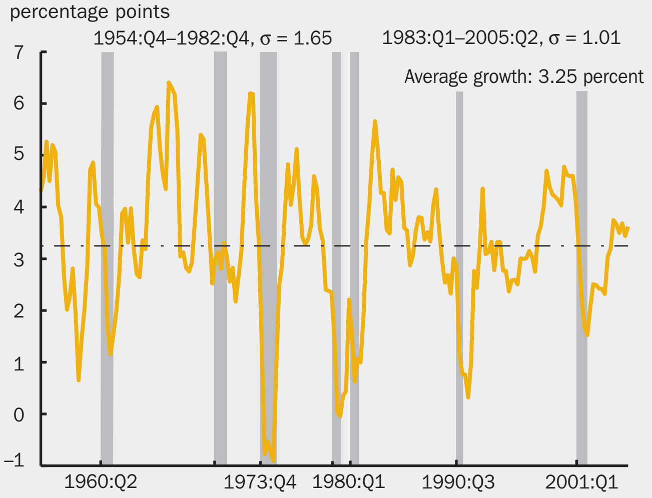

Notes: The figure plots the four-quarter growth rate of personal consumption expenditures on nondurable goods and nonhousing services. The shaded areas are recessions as identified by the National Bureau of Economic Research. See the text for further details.

Data

We must first address an important measurement issue before taking our model to the data. The distinction between durable goods and services is subtle. A system of national income accounts can either treat a durable goods purchase as an investment (recorded as an expenditure by the business sector) that yields a flow of rental services to households or as a direct expenditure by the household sector. The U.S. NIPA do not treat this issue consistently. The NIPA treat housing as a business transaction: The construction of a home is recorded as residential investment, and the flow of services from the housing stock is recorded as consumption expenditures. In contrast, the purchase of any other new household durable good is counted only as a consumption expenditure: The service flow from the durable does not appear anywhere in the national accounts. For the purposes of this article, we assume that the household sector purchases all durables directly. In practice, this means that we subtract the service flow from housing from the NIPA consumption data and from GDP when constructing our measures of nondurables and services consumption and total private income.[7] We also measure household expenditures on durables as the sum of expenditures on durables and residential investment.

Next, it is helpful to examine the histories of the five variables of interest. We begin with the model's core variable, consumption of nondurables and services. Figure 1 plots $\Delta c_t = \ln(C_t / C_{t-1})$. For growth rates sufficiently close to zero, this approximately equals $(C_t - C_{t-1})/C_t$. Its mean is 3.25 percent. This average growth rate does not drift over time, which indicates that there are no long-run changes in the growth rate of consumption that could confound our analysis.[8] However, the volatility of Δc_t diminished substantially in the 1980s. Its standard deviation over the 1954–82 period was 1.65 percent; since then, the standard deviation has been 1 percent. McConnell and Perez-Quiros (2000) documented that the variances of most NIPA expenditure categories declined at about this time.

To impose long-run balanced growth on the model's forecasts, we work with the ratios of expenditures on durable goods and nominal privately demanded GDP to nominal nondurable consumption. Figures 2 and 3 plot the logarithms of these ratios from 1954:Q1 through 2005:Q2. According to the balanced growth hypothesis, neither ratio should persistently drift away from its mean. This was the case through the early 1980s, but since then both of these series have persistently declined. One way to measure the drift in these data is shown in the dashed lines, which plot the averages of the series' previous 40 quarters' values. These moving averages are stable in the 1960s and 1970s, trend

FIGURE 2

Ratio of durable to nondurable consumption expenditures

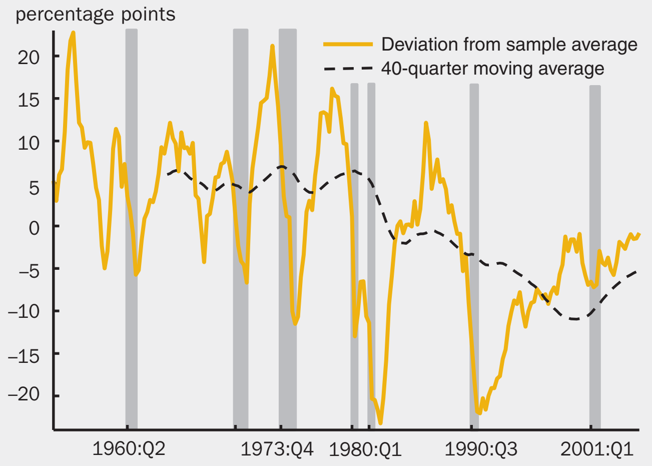

Notes: The figure plots the logarithm of the ratio of nominal expenditures on durable goods and residential investment to nominal expenditures on nondurable goods and nonhousing services. The dashed line is a 40-quarter one-sided moving average of this ratio, and the shaded areas are recessions as identified by the National Bureau of Economic Research. See the text for further details.

FIGURE 3

Ratio of private GDP to nondurable consumption expenditures

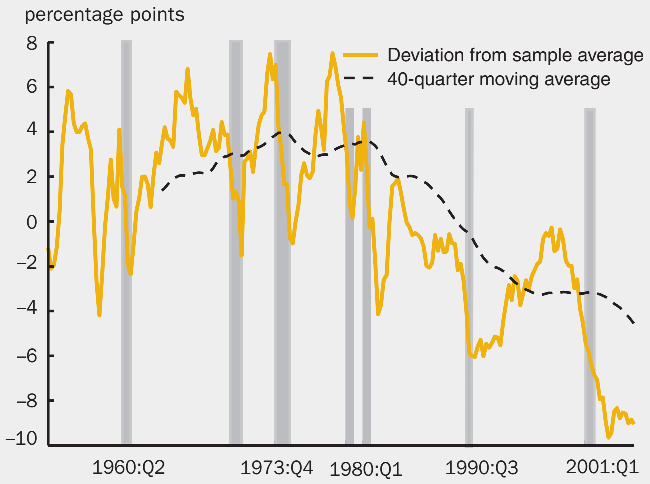

Notes: The figure plots the logarithm of the ratio of nominal gross domestic product (GDP) less government expenditures and housing services to nominal expenditures on nondurable goods and nonhousing services. The dashed line is a 40-quarter one-sided moving average of this ratio, and the shaded areas are recessions as identified by the National Bureau of Economic Research. See the text for further details.

down in the 1980s, and stabilize to some degree in the 1990s. The shift down in the values of the private income ratio reflects the well-known decline in the U.S. savings rate.

Since nondurable consumption is our fundamental indicator of the income process, we measure inflation using the corresponding implicit price index. Panel A of figure 4 plots this index's inflation rate. Over the post-Korean War sample, inflation underwent some dramatic changes. It rose in the late 1960s, surged with the oil shocks in the 1970s, and then fell in the 1980s and 1990s. The inflation movements over the past 25 years, however, are small relative to the swings that occurred between the late 1960s and early 1980s. Panel B of figure 4 plots the nominal interest rate on federal funds. Its changes track inflation's well, but it declined from its peak in the early 1980s somewhat more slowly.

Sample period

The decline of macroeconomic volatility and persistent movements in nominal expenditure ratios pose problems for an econometric model based on stable relationships. Recent research has documented and tried to explain the marked changes in a wide range of many macroeconomic aggregates that apparently have occurred since the early 1980s.[9] Campbell and Hercowitz (2004) argue that the decline in macroeconomic volatility reflects in part greater household access to credit markets. Other structural changes that we expect to influence household decisions in the long run include the elimination of consumer credit interest deductions and the development of secondary markets for mortgage debt. Structural changes in production and distribution may also be a factor; for example, Kahn, McConnell, and Perez-Quiros (2002) argue that improved inventory management technology has helped reduce the volatility of GDP. Monetary policy has also been a source of structural change. The rise in inflation that began in the late 1960s reflected in part economic policies that favored economic growth over inflation stabilization, while Paul Volcker's increased attention to inflation in 1979 began the transition to the current era of very low inflation.

Explaining structural changes such as these lies outside of the scope of our forecasting model. Nevertheless,

FIGURE 4

Inflation and the federal funds rate

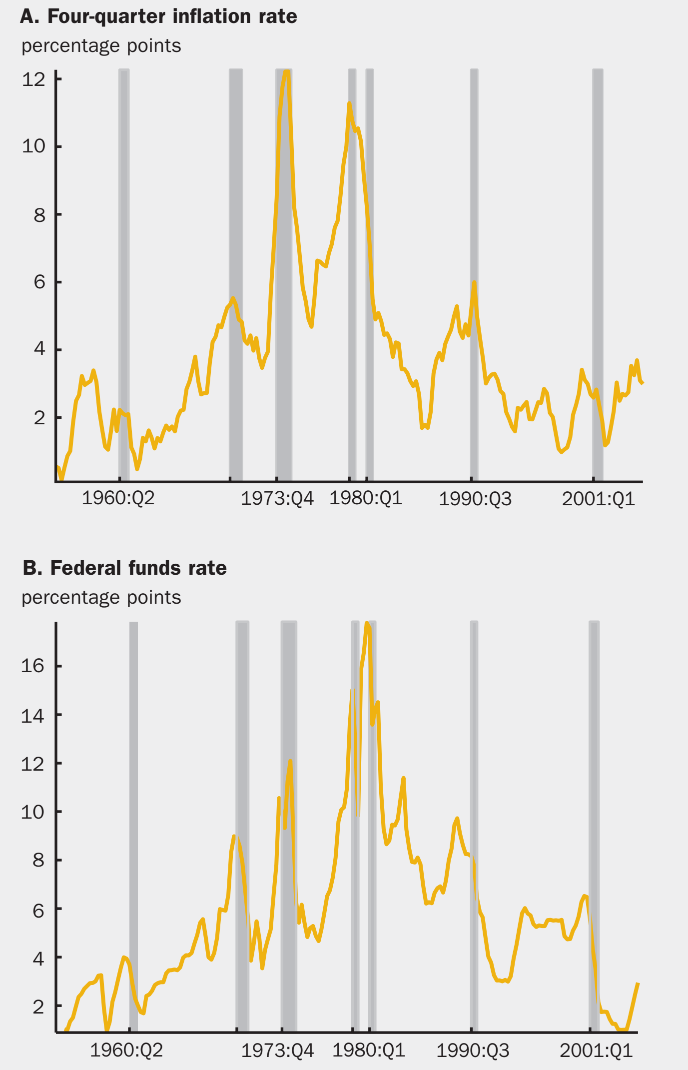

Notes: Inflation is measured using the chain-weighted price deflator for personal consumption expenditures for nondurable goods and nonhousing services. Both variables are measured at annual rates. The shaded areas are recessions as identified by the National Bureau of Economic Research. See the text for further details.

their presence does not eliminate the value of imposing balanced growth on our forecasts because after some adjustment period, the nominal ratios will settle down to their new balanced growth path. Accordingly, we decided to make two adjustments to our model. First, we forecast the deviations of the two nominal ratios from their 40-quarter moving averages instead of the nominal ratios themselves. This procedure removes the drift in these ratios' means arising from structural change.[10] Second, we use only data from 1983:Q1 through the present. We do so because the change in monetary policy regime and factors that may have reduced economic volatility also plausibly changed the way economic agents respond to a variety of economic shocks.

This sample period is much shorter than those used to estimate most macroeconomic forecasting models, which generally start in the late 1960s or even earlier. This greatly limits the number of parameters that we can estimate with sufficient statistical precision. This need for parsimony dictated by the use of a short sample complements the balanced-growth arguments for considering a model that forecasts only a few broad NIPA aggregates. It also leads us to use exclusion tests to further restrict the number of model parameters.

The stochastic simultaneous equations model

The first step in the construction of our forecasts is the estimation of a small structural model of the U.S. economy. The model consists of a system of simultaneous equations for the five variables of interest. Because the equations allow a number of shocks to simultaneously influence multiple variables, we refer to this system as the stochastic simultaneous equations model. This model imposes restrictions from theory in order to identify permanent and transitory shocks to income. In this sense, the model is *structural*.

The model's first equation is for the growth rate in the consumption of nondurables and services. This equation also defines the permanent income shock, ε_t.

$$8)\quad \Delta c_t = \mu_c + \sum_{l=1}^{m_1} \alpha_{lc} \Delta c_{t-l} + \sum_{l=0}^{m_1} \alpha_{l\pi} \pi_{t-l} + \sum_{l=0}^{m_1} \alpha_{lr} r_{t-l} + \varepsilon_t,$$

where π_t and r_t denote the annualized one-quarter inflation rate and the federal funds rate. This specification embodies the theoretical restriction that only consumption and interest rates predict future consumption. However, the presence of π_t and r_t on its right-hand side provides a channel for transitory shocks to immediately impact Δc_t.

BOX 1

The model's long-run restrictions

The restriction that only the permanent income shock can permanently change the level of nondurable consumption implies a set of restrictions on the unknown coefficients in equation 8. We derive these restrictions given the assumption that $\sum_{l=1}^{m_1}\alpha_{lc} < 1$. This implies that no shock permanently changes the *growth rate* of nondurable consumption. This is a prediction of most general-equilibrium business cycle models.

To begin, suppose that a particular transitory shock occurs in quarter t that impacts both inflation and interest rates. Denote the changes in the *level* of nondurable consumption, inflation, and the interest rate in quarter $t+j$ due to the shock with ∇c_{t+j}, $\nabla \pi_{t+j}$ and ∇r_{t+j}. We assume that the long-run responses of inflation and interest rates to the shock both equal zero so that $\nabla\pi = \lim_{j\to\infty}\nabla\pi_{t+j} = 0$ and $\nabla r = \lim_{j\to\infty}\nabla r_{t+j} = 0$. We wish to characterize $\nabla c = \lim_{j\to\infty}\nabla c_{t+j}$ so that we can find the restrictions required to set it equal to zero. From equation 8 we can write ∇c_{t+j} as

$$\nabla c_{t+j} = \nabla c_{t+j-1} + \sum_{l=1}^{m_1}\alpha_{lc}\left(\nabla c_{t+j-l} - \nabla c_{t+j-l-1}\right) + \sum_{l=0}^{m_1}\alpha_{l\pi}\nabla\pi_{t+j-l} + \sum_{l=0}^{m_1}\alpha_{lr}\nabla r_{t+j-l}.$$

Summing these equations for $j = 0,1,\ldots, M$ yields

$$\nabla c_{t+M} = \sum_{l=1}^{m_1}\alpha_{lc}\nabla c_{t+M-l} + \left(\sum_{l=0}^{m_1}\alpha_{l\pi}\right)\sum_{j=0}^{M}\nabla\pi_{t+j-l} - \sum_{l=1}^{m_1}\alpha_{l\pi}\sum_{j=0}^{l-1}\nabla\pi_{t+M-j} + \left(\sum_{l=0}^{m_1}\alpha_{lr}\right)\sum_{j=0}^{M}\nabla r_{t+j-l} - \sum_{l=1}^{m_1}\alpha_{lr}\sum_{j=0}^{l-1}\nabla r_{t+M-j}.$$

As M becomes very large, ∇c_{t+M} and its m_1 lagged values all approach ∇c. Because both $\nabla\pi$ and ∇r equal zero, it satisfies

$$\left(1 - \sum_{l=1}^{m_1}\alpha_{lc}\right)\nabla c = \left(\sum_{l=0}^{m_1}\alpha_{l\pi}\right)\sum_{j=0}^{\infty}\nabla\pi_{t+j} + \left(\sum_{l=0}^{m_1}\alpha_{lr}\right)\sum_{j=0}^{\infty}\nabla r_{t+j}.$$

Neither $\sum_{j=0}^{\infty}\nabla\pi_{t+j}$ nor $\sum_{j=0}^{\infty}\nabla r_{t+j}$ necessarily equals zero, so we require both $\sum_{l=0}^{m_1}\alpha_{l\pi}$ and $\sum_{l=0}^{m_1}\alpha_{lr}$ to equal zero for ∇c to equal zero. These are the two restrictions we impose on the model.

The equation defines ε_t as the permanent income shock by restricting it to be the only shock that has a long-run impact on c_t. In practice, ε_t encompasses many fundamentals such as changes in technology, regulation, and access to financial markets. Gathered together, these are all of the factors causing permanent and equal-sized changes in C_t, X_t, and Y. Other shocks may have transitory effects, but none have any permanent ones.

As a technical matter, as shown in box 1, the two necessary and sufficient conditions for ε_t to be the only shock that has a long-run impact on c_t are $\sum_{l=0}^{m_1}\alpha_{l\pi} = 0$ and $\sum_{l=0}^{m_1}\alpha_{lr} = 0$. These conditions guarantee that any negative influence of a shock to either interest rates or inflation on nondurable consumption growth in the current quarter is offset by a positive influence in the following quarters, so the net effect on c_t is zero. We impose these restrictions on our estimates of equation 8's unknown parameters.

The model's remaining equations allow ε_t to immediately impact all variables; in particular, ε_t may contemporaneously influence the values of π_t and r_t that enter equation 8. This implies that the equation's error term and its right-hand side variables are correlated, so we cannot estimate its unknown parameters using ordinary least squares regression. Instead, we assume that there does not exist any information available before time t that is useful in forecasting ε_t. This is a natural restriction to apply to an economic shock. Mathematically, this means that the expectation of ε_t is independent of the lagged values of the data in the model; that is:

$$\text{9)}\quad E[\varepsilon_t] = E[\varepsilon_t\Delta c_{t-j}] = E[\varepsilon_t\pi_{t-j}] = E[\varepsilon_t r_{t-j}] = E[\varepsilon_t y_{t-j}] = E[\varepsilon_t x_{t-j}] = 0$$

for all $j \geq 1$. Here, y_t and x_t are the deviations from 40-quarter moving averages in the logarithms of the nominal ratios of private GDP and durable goods purchases to nondurable consumption. The conditions in equation 9 suggest estimating μ_c and the α_{ij}'s by using the values that set the sample covariances between ε_t and the $t-j$ dated data equal to zero. However, there are more covariances to set equal to zero than there are coefficients to choose. This leads us to use linear combinations of the covariances—which also have an expected value equal to zero—to estimate the equation. We use Hansen's (1982) generalized method of moments (GMM) estimator, which identifies the particular linear combinations that yield the most precise estimates.[11]

The model's second equation is:

$$10)\quad \pi_t = \mu_\pi + \sum_{l=0}^{m_2}\beta_{lc}\Delta c_{t-l} + \sum_{l=1}^{m_2}\beta_{l\pi}\pi_{t-l} + \sum_{l=0}^{m_2}\beta_{lr}r_{t-l} + \sum_{l=0}^{m_2}\beta_{ly}y_{t-l} + \sum_{l=0}^{m_2}\beta_{lx}x_{t-l} + e_{\pi t}.$$

Because π_t appears on its left-hand side, we call this the inflation equation. The error term in equation 10 has no concrete economic interpretation. Instead, we treat it as a statistical forecast error. Because it is a forecast error, its covariance with any variable that could be used to forecast it equal zero. That is,

$$E[e_{\pi t}] = E[e_{\pi t}\Delta c_{t-j}] = E[e_{\pi t}\pi_{t-j}] = E[e_{\pi t}r_{t-j}] = E[e_{\pi t}y_{t-j}] = E[e_{\pi t}x_{t-j}] = 0$$

for all $j \geq 1$. Thus we can use the corresponding sample covariances to estimates the β's. In addition, because we restrict ε_t to be the *only* shock with a long-run impact on c_t, we also impose the restriction that $e_{\pi t}$ is independent of ε_t. Mathematically, this means $E[e_{\pi t}\varepsilon_t] = 0$, which is one more covariance to use in estimation. Just as with equation 9, we use the GMM estimator to estimate its unknown parameters.

The specification and estimation of the remaining three forecasting equations proceed similarly. In each of them, r_t, y_t, and x_t take the role of π_t in equation 10. We selected the lag lengths ($m_1,\ldots,m_5$) for each equation as the smallest value of m such that the coefficients multiplying an additional lag's variables are jointly statistically insignificant.

VAR results

Table 1 displays the estimated system of simultaneous equations using matrix notation. The system is parsimonious: Only the inflation equation has more than two lags. The coefficient multiplying r_t in equation 8 and the coefficients multiplying Δc_t in the inflation and interest rate equations are statistically significant. This implies that transitory shocks can affect consumption growth immediately and that the permanent income shock can affect both inflation and interest rates on impact. However, the coefficients multiplying Δc_t in the two expenditure-ratio equations are not statistically significant. Thus, the initial changes in nondurable consumption are proportional to the changes in income and durable consumption expenditures.

The apparent ability of transitory economic shocks to influence the growth rate of nondurable consumption sharply contrasts with Hall's (1978) theoretical prediction that all changes in consumption are permanent. However, Hall's result depends on the assumption of a constant interest rate. In more general models with a market-determined interest rate, shocks that *temporarily* change the economy's productive capability will temporarily affect both the level of consumption and the interest rate. The estimated impact of transitory shocks on consumption growth in our model demonstrates this possibility's empirical relevance.

Before proceeding to use the model for forecasting, we wish to examine whether it displays evidence of specification error. To do so, we created figure 5, which plots the actual values of the model's variables along with the estimated model's forecasts for them given the data known in 1983:Q4.[12] Structural change in the sample would cause the predicted values to drift away from the actual data. In fact, the data track the initial forecasts well.

Impulse responses to the permanent income shock

The permanent income theory makes very specific predictions for the responses of Δc_t and x_t to a permanent income shock. Figure 6 plots the estimated model's counterparts to these. Each panel plots the response of a variable of interest over time to a single positive permanent income shock equal in magnitude to the shock's estimated standard deviation. As the theoretical model suggests, consumption of nondurables and services reacts significantly and quickly, it increases about 0.27 percent when the shock hits and achieves its complete increase of 0.44 percent after three quarters.

The adjustment of household capital lasts longer. On impact, X_t rises about 0.60 percent. The response increases to 1.36 percent after two quarters, remains at this level for about another year, and then slowly moves down toward the same permanent increase found in C_t. This pattern demonstrates how expenditures on durable goods must temporarily rise more than C_t in order to bring the stock of durable goods into balance with consumption. However, the persistence of the increased expenditures is quite protracted relative to the theoretical model's predicted transitory increase. We believe that the drawn-out response reflects costs of quickly adjusting the stock of durable goods that are absent from the model.

TABLE 1

The estimated structural VAR

$$
\begin{bmatrix}
1.00 & -0.08 & -0.66^{***} & 0.00 & 0.00 \\
2.74^{***} & 1.00 & 0.00 & 0.00 & 0.00 \\
0.56^{***} & 0.00 & 1.00 & 0.00 & 0.00 \\
0.20 & 0.00 & 0.00 & 1.00 & 0.00 \\
-1.08 & 0.00 & 0.00 & 0.00 & 1.00
\end{bmatrix}
\begin{bmatrix} \Delta c_t \\ \pi_t \\ r_t \\ y_t \\ x_t \end{bmatrix}
=
\begin{bmatrix}
-0.18 & -0.10^{**} & -1.04^{***} & 0.00 & 0.00 \\
1.28^{***} & 0.19 & -0.74^{*} & 0.34 & -0.10 \\
0.54^{***} & 0.12^{***} & 1.20^{***} & 0.31^{***} & 0.05^{*} \\
0.16 & 0.07 & 0.27^{*} & 0.91^{***} & 0.13^{***} \\
0.71 & -0.56^{***} & -0.00 & -0.06 & 0.93^{***}
\end{bmatrix}
\begin{bmatrix} \Delta c_{t-1} \\ \pi_{t-1} \\ r_{t-1} \\ y_{t-1} \\ x_{t-1} \end{bmatrix}
$$

$$
+
\begin{bmatrix}
0.22 & 0.03 & 0.38^{***} & 0.00 & 0.00 \\
0.85^{**} & 0.32^{**} & 0.79 & 0.13 & 0.38^{***} \\
-0.17 & 0.04 & -0.29^{**} & -0.25^{***} & -0.03 \\
-0.05 & -0.01 & -0.29^{*} & 0.06 & -0.16^{***} \\
0.00 & 0.00 & 0.00 & 0.00 & 0.00
\end{bmatrix}
\begin{bmatrix} \Delta c_{t-2} \\ \pi_{-2} \\ r_{t-2} \\ y_{t-2} \\ x_{t-2} \end{bmatrix}
+
\begin{bmatrix}
0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\
0.23 & 0.22^{*} & 0.16 & -0.59^{*} & -0.16 \\
0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\
0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\
0.00 & 0.00 & 0.00 & 0.00 & 0.00
\end{bmatrix}
\begin{bmatrix} \Delta c_{t-3} \\ \pi_{t-3} \\ r_{t-3} \\ y_{t-3} \\ x_{t-3} \end{bmatrix}
$$

$$
+
\begin{bmatrix}
0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\
-1.40^{***} & 0.04 & -0.18 & 0.07 & -0.04 \\
0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\
0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\
0.00 & 0.00 & 0.00 & 0.00 & 0.00
\end{bmatrix}
\begin{bmatrix} \Delta c_{t-4} \\ \pi_{t-4} \\ r_{t-4} \\ y_{t-4} \\ x_{t-4} \end{bmatrix}
+
\begin{bmatrix}
\text{Permanent Income Shock}_t \\
\text{First Unidentified Transitory Shock}_t \\
\text{Second Unidentified Transitory Shock}_t \\
\text{Third Unidentified Transitory Shock}_t \\
\text{Fourth Unidentified Transitory Shock}_t
\end{bmatrix}
$$

*Significant at the 10 percent level.
**Significant at the 5 percent level.
***Significant at the 1 percent level.

Notes: Δc_t refers to the real logarithmic growth rate of personal consumption expenditures on nondurable goods and nonhousing services, π_t refers to inflation rate measured using that series' chained price deflator, r_t denotes the nominal interest rate on federal funds, and y_t and x_t refer to the logarithmic ratios of nominal private gross domestic product (GDP) and nominal durable goods expenditures to nominal nondurable consumption, adjusted as described in the text. See the text for further details.

The impulse responses of privately demanded GDP also exhibit a drawn-out and cyclical pattern. Initially, Y_t rises by 0.22 percent; it then climbs to 0.41 percent after two quarters. Instead of remaining at this level—which is very close to its long-run response—Y_t then falls over the next five years to approximately 0.25 percent above its initial level. Thereafter, it begins a slow climb and achieves its long-run level 12 years after the initial shock. It appears that the permanent increase in consumption substantially precedes the realization of additional income. The permanent income theory emphasizes household borrowing to increase consumption following increases in future wealth. These estimated responses indicate that this pattern has empirical relevance for the U.S. economy. However, the borrowing is surprisingly persistent. We feel that this pattern warrants further investigation.

The responses of inflation and interest rates shed some light on the monetary response to a permanent income shock. On impact, the shock induces inflation to fall by 0.74 percent. The initial disinflation quickly reverses itself. After four quarters, inflation is about 0.10 percent above its initial value. Thereafter, the inflation rate tends to exceed its initial level for a considerable period and return to its pre-shock level at a measured pace. Although the initial response is deflationary, the estimated long-run response of the price *level* is a 1 percent increase. The estimated response of interest rates sheds some light on this response. The permanent income shock induces a drop of 16 basis points in the federal funds rate, which lasts nearly one year. In the second year after the shock, the interest

FIGURE 5

The model's initial forecasts

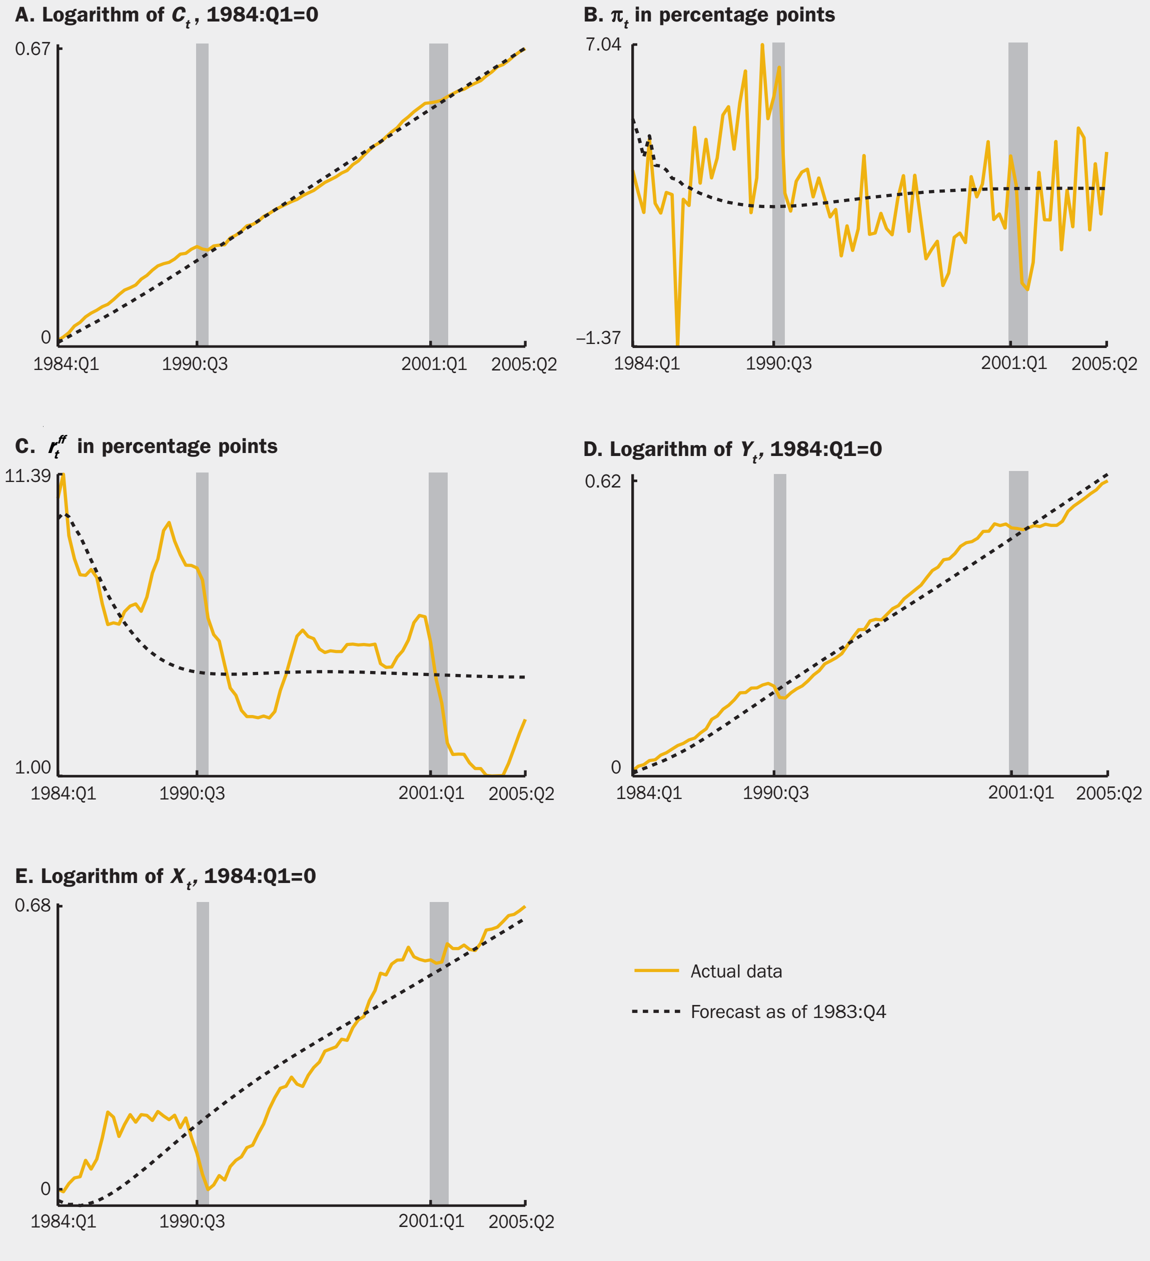

Notes: The shaded areas are recessions as identified by the National Bureau of Economic Research. See the text for further details.

rate rises towards its original level, but remains about 7 basis points below it. In light of the initial deflation, we can interpret the drop in the interest rate as an accommodative policy response to the apparent absence of inflationary pressure.

Variance decompositions

How important are permanent income shocks for economic fluctuations? Table 2 reports the fractions of forecast uncertainty due to the permanent income shock and due to the transitory shocks taken together at various forecasting horizons. For consumption, the permanent income shock accounts for a large

FIGURE 6

Impulse-response functions from the estimated model

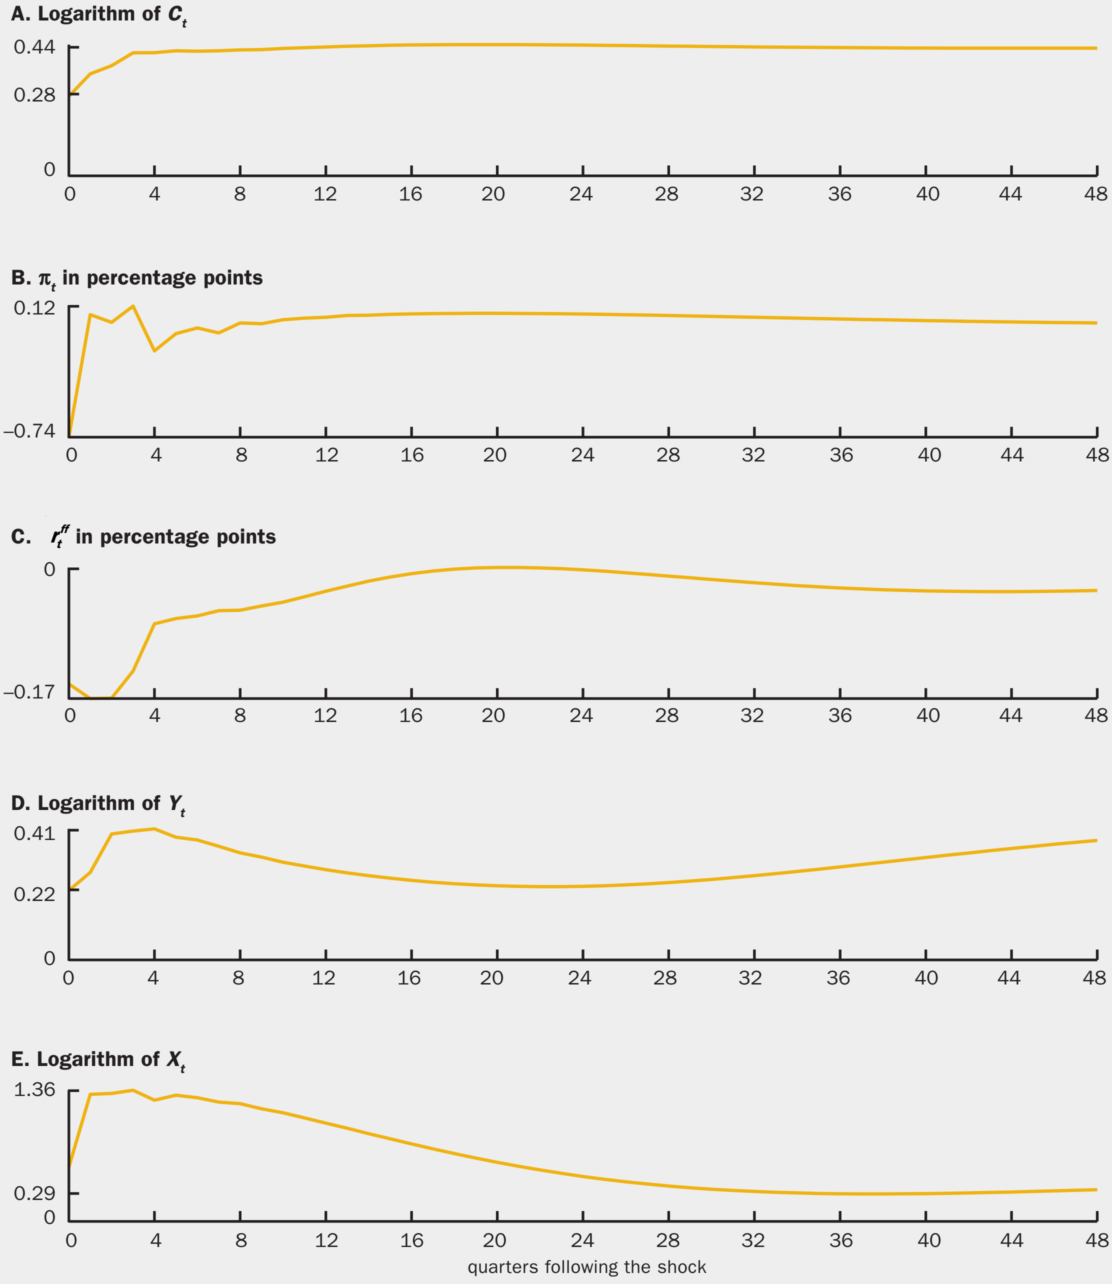

Notes: The figure plots responses to a one standard deviation shock. See the text for further details.

percentage of the movement at all forecast horizons; it explains 63 percent of the variance of the one-quarter-ahead forecast error, 73 percent of the four-quarter-ahead error's variance, and nearly all of the five-year-ahead error. Still, many temporary factors influence the economy, and these explain a substantial fraction of nondurable consumption's forecast-error variance over the first year or two following the shock.

These results highlight an important difference between our model and Cochrane's (1994). He identifies the permanent shock by restricting it to be the only factor affecting consumption in the current period. So, by construction, his transitory shock explains

TABLE 2

The model's variance decompositions

Shock to	Variance of c_t	π_t	r_t	y_t	x_t
	One quarter ahead				
Permanent income	62.6	40.2	17.9	12.3	8.8
Others	37.4	59.8	82.1	87.7	91.2
	Four quarters ahead				
Permanent income	73.1	29.3	7.8	16.2	27.4
Others	26.9	70.7	92.2	83.8	72.6
	Eight quarters ahead				
Permanent income	77.8	28.2	3.7	19.4	34.6
Others	22.2	71.8	96.3	80.6	65.4
	20 quarters ahead				
Permanent income	88.1	27.7	2.5	23.9	44.5
Others	11.9	72.3	97.5	76.1	55.5

none of the one-quarter-ahead forecast error in consumption. After about ten quarters, however, this shock accounts for about one quarter of the forecast error in consumption, and it remains at this level of importance for all subsequent forecast horizons. In contrast, our identifying assumptions allow transitory factors to influence the near-term forecast, but restrict them from having any influence on the long-run outlook for consumption.

Next, consider the forecast-error variances for the other spending variables in our model. Transitory factors explain most of the forecast errors in durable goods expenditures and total private income over the short and medium terms. For example, even at the two-year horizon, the transitory shocks explain 80 percent of private GDP's forecast-error variance. By construction, the influence of the transitory shocks falls to zero as the forecast horizon increases. However, this takes a long time to occur. At a 20-year horizon, the transitory shocks still account for 36 percent of the forecast error in Y_t and 56 percent of the error in X_t.[13]

The permanent shock has an important influence on inflation and interest rates in the very near term. It accounts for about 40 percent of the one-quarter-ahead forecast error in inflation and nearly 20 percent of that for the interest rate. Its influence on the interest rate drops quickly. At the two-year horizon, it accounts for only 3.7 percent of the forecast-error variance. This is consistent with economic intuition. As we noted previously, the long-run equilibrium real interest rate depends on the trend in the growth rate of consumption but does not change with shocks to the level of output. The influence of permanent income shocks on inflation variance also falls as the forecast horizon increases, but not as rapidly.

What does the permanent income shock say about history?

During the 1983–2005 period that we use to estimate our model, the U.S. economy experienced two recessions, a productivity boom, and a decline in inflation. Does the shock to permanent income that we estimate have anything interesting to say about these developments? Consider figure 7. The solid line in each panel plots the deviations of the variable from the forecast based on the data in hand in 1983:Q4 (the forecasts shown in figure 5, p. 62) and the dashed line plots the path of this deviation if *only* the permanent income shock occurred.

The lines for nondurable consumption match closely, highlighting the importance of the permanent income shock in explaining consumption. The model attributes much of the weakness in consumption during the 1990–91 recession to a decline in permanent income. It also identifies permanent factors as important contributors to spending during the boom in the second half of the 1990s. In contrast, the model estimates that most of the modest decline in consumption growth during the 2001 recession can be explained by transitory factors. Put differently, the model interprets the 1990–91 recession as being a more serious decline of the economy's permanent productive capacity than the most recent economic downturn.

The role of permanent income shocks in the business cycle fluctuations of x_t and y_t is qualitatively similar: It explains some of the movements in spending on durables and total income during the 1990–91 recession and 1995–2001 boom, but is not as important in explaining the last recession. Quantitatively, however, we see that a much larger portion of the declines in x_t and y_t during the last two recessions were due to transitory components. Furthermore, much of the run-up in income during the second half of the 1990s was transitory. Recall that business investment boomed during the period, and the model interprets much of that increase as reflecting transitory factors.

FIGURE 7

Permanent income shocks' contribution to variance

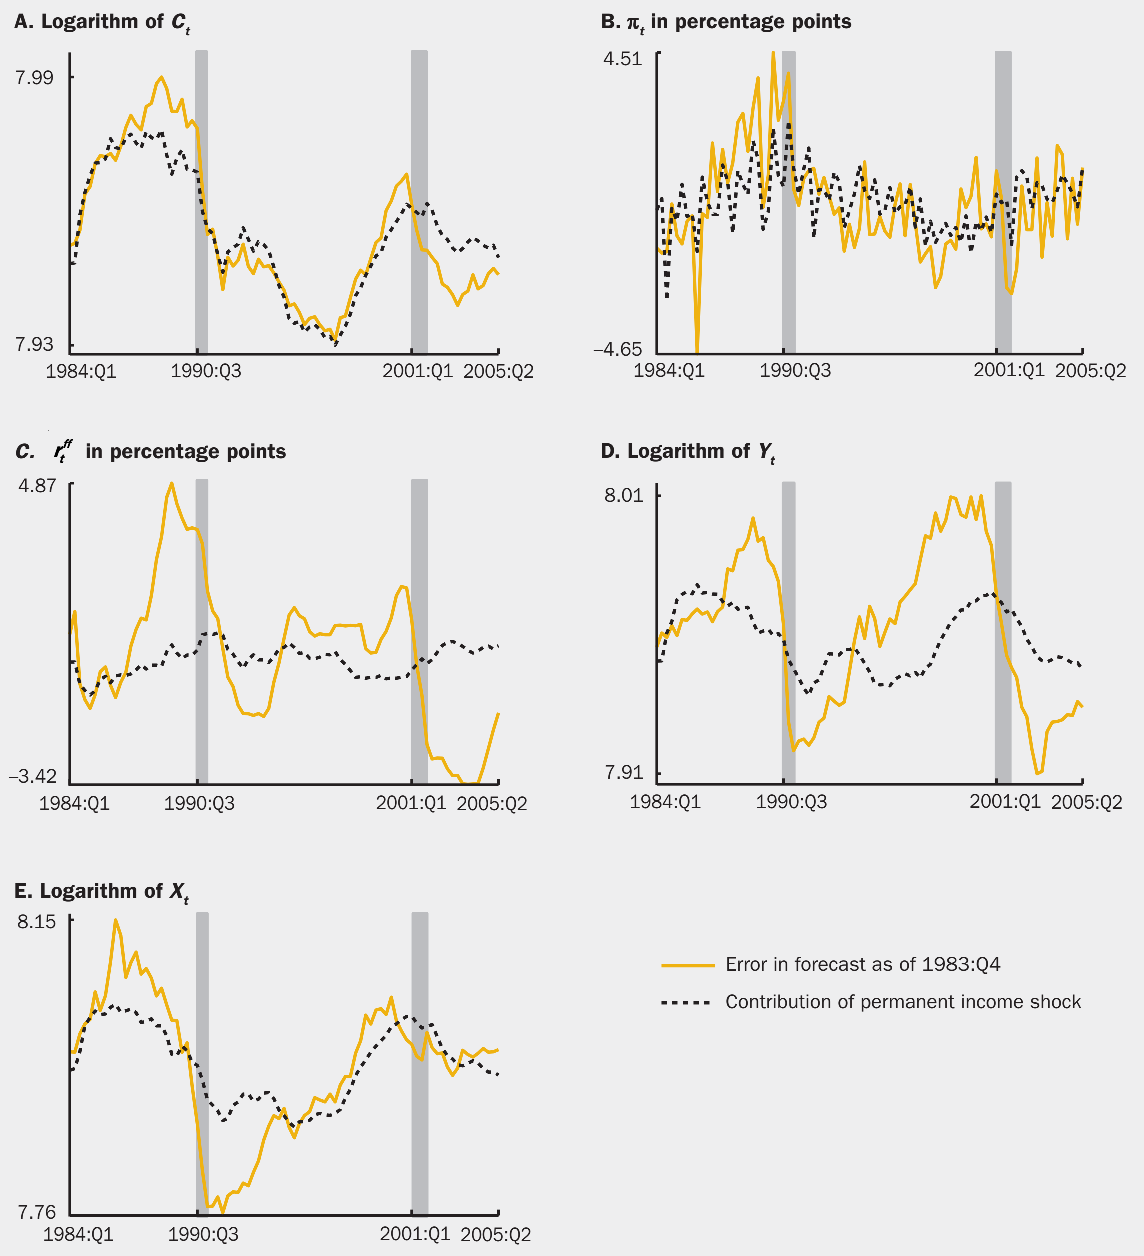

Notes: The shaded areas are recessions as identified by the National Bureau of Economic Research. See the text for further details.

Forecasts

With the model's estimates in hand, constructing forecasts of its variables is relatively straightforward. Box 2 contains a detailed description of the process. Table 3 reports the four-quarter growth rates for 2003 and 2004 and their model-based forecasts for 2005 and 2006 for nondurable consumption expenditures, privately demanded GDP, durable consumption expenditures, and the price level. These forecasts use the information available as of August 1, 2005. Recall that the model's definitions of nondurable consumption, inflation, private GDP, and durable goods purchases are nonstandard, so the data and forecasts of table 3 are not directly comparable to the identically labeled values reported in the NIPA by the U.S. Bureau of Economic Analysis.

The model's long-run consumption growth rate is 3.29 percent. Growth of nondurable consumption expenditures fell slightly short of this in 2003 and exceeded it by 0.40 percentage point in 2004. The forecasts of nondurable consumption growth for both 2005 and 2006 differ little from the long-run growth rate. By construction, the model's long-run growth rate of output equals that for consumption. In fact, output growth exceeded this level by 0.50 percentage point in 2003 and fell 0.30 short of it in 2004. The model-based forecast of private GDP growth for 2005 is 0.20 percentage point below the long-run growth rate, while that for 2006 falls short by a more substantial 0.50 percentage point. The growth rate of durable goods purchases is projected to slow markedly. Given that our measure of durable goods purchases includes residential investment, this is not surprising because the growth rate of X_t substantially exceeded the model's long-run growth rate in both 2003 and 2004. Table 3's last row reports realized annual inflation rates and their forecasts. The model's long-run inflation rate is 3 percent. Inflation fell below this in 2003 by 0.30 percentage point, and exceeded it by about 0.70 percentage point in 2004. The model forecasts that inflation will exceed its average in both 2005 and 2006, but by no more than 0.30 percentage point.

The output gap

Traditional approaches to monetary policy depend on comparing the pace of economic activity to its potential. In practice, the definition of potential output growth is problematic and its measurement correspondingly difficult.[14] Indeed, these difficulties have led Hall (2005) to advocate dispensing with the concept of potential output altogether. The structural model we present suggests one approach to this problem. We define "potential" output as the value of private GDP if the only variables affecting it were the permanent income shocks and the structural changes taken out by the 40-quarter moving averages of the two nominal ratios. The difference between this and the actual level of private GDP is our measure of the output gap. A full exploration of

BOX 2

Generating forecasts from the model

Generating forecasts from the model is a two-step process. We start with the past four quarters' values of the model's variables, Δc_t, π_t, r_t, y_t, and x_t. This information allows us to calculate the right-hand sides of the model's equations. The current quarter's values of these variables appear in the equations' left-hand sides. We use back substitution to solve these equations and calculate our forecasts of the current quarter's values. We denote these with $\Delta c_{t|t-1}$, $\pi_{t|t-1}$, $r_{t|t-1}$, $y_{t|t-1}$, and $x_{t|t-1}$ to explicitly mark their dependence on the information available at the end of quarter $t-1$. We then repeat this procedure to calculate $\Delta c_{t+1|t-1}$, $\pi_{t+1|t-1}$,… using $\Delta c_{t|t-1}$, $\pi_{t|t-1}$,… to replace the unknown values of the model's variables in period t. Repeating this process generates forecasts for all of the model's variables for any desired horizon.

This first step directly generates our forecasts for nondurable consumption growth and inflation. Because the model forecasts *deviations* of the ratios of private GDP and durable consumption expenditures to nondurable consumption from their 40-quarter moving averages, we must adjust the forecasts generated by the first step. Consider first $x_{t|t-1}$. Because the 40-quarter moving average of this nominal ratio's past values is available, we can account for its subtraction from the original data by adding it back to $x_{t|t-1}$. That is, if we denote the original unadjusted value of the nominal ratio with z_t, then

$$z_{t|t-1} = x_{t|t-1} + \frac{1}{40}\sum_{j=1}^{40} z_{t-j}.$$

To adjust $x_{t+1|t-1}$, we use this forecast and the past 39 values of the nominal ratio to forecast the 40-quarter moving average applicable in quarter $t+1$ and add the result to $x_{t+1|t-1}$. So,

$$z_{t+1|t-1} = x_{t+1|t-1} + \frac{1}{40}\left(z_{t|t-1} + \sum_{j=1}^{39} z_{t-j}\right).$$

Adjusting the forecasts of this nominal ratio at later horizons proceeds analogously. Transforming these adjusted forecasts and the forecasts of nondurable consumption growth into forecasts of durable consumption expenditures is straightforward. The procedure for adjusting $y_{t+j|t-1}$ and forecasting private GDP is the same.

TABLE 3

The model's forecasts of annual growth rates as of 2005:Q2

	2003:Q4/ 2002:Q4	2004:Q4/ 2003:Q4	2005:Q4/ 2004:Q4	2006:Q4/ 2005:Q4
C_t	3.20	3.68	3.25	3.35
Y_t	3.79	3.00	3.11	2.77
X_t	6.14	4.38	3.36	3.48
P_t	2.65	3.69	3.14	3.28

Notes: C_t, Y_t, X_t, and P_t refer to nondurable consumption, private gross domestic product, durable goods purchases, and the price level. See the text for further details.

this measure's relationship with inflation and monetary policy lies well beyond the scope of this present article. Here, we restrict ourselves to a simple discussion of its evolution over past and current business cycles.

One can roughly gauge the path of the output gap by examining panel D of figure 7 (p. 65). Our measure of the output gap equals the vertical distance between the solid and dashed lines. Apparently, the output gap equaled zero at both of the sample's business cycle peaks as defined by the National Bureau of Economic Research. Following the 1990 peak, the gap remained negative for about three years, underscoring the impression of a lengthy recovery from that recession. Similarly, the gap has remained negative since the 2001 peak.

An advantage of using a structural model for our forecasts is that constructing forecasts of this definition of the output gap is relatively straightforward. Figure 8 plots the actual value of the gap since 1999 as well as the model-based forecasts through the end of 2006. At the beginning of 1999, transitory shocks had driven output above what it would otherwise be by 4 percentage points. The third quarter of 2000 began a precipitous drop in the output gap. It achieved its recent trough value of –3.8 percent in 2003:Q1. The gap finished our sample period at –1.2 percent. The forecast of the gap shows it rising through the end of 2006, but not by much. Its forecasted value for 2006:Q4 is –0.8 percent. By construction, the model's forecasts of the gap approach zero as the forecast horizon increases. However, this convergence takes approximately 20 years to complete. Thus, the model implies that returning the output gap to zero or above in the near term would require some type of favorable transitory shock.

Conclusion

This article constructs a small macroeconomic forecasting model of real economic activity and inflation. The model identifies permanent and transitory shocks to output using only a few simple assumptions. First, the theory of permanent income tells us that real consumption of nondurables and services is uniquely informative about the long-run income prospects of the economy. Second, in the U.S. economy nominal outlays on durable goods and total private income do not persistently drift away from spending on nondurable consumption. These identifying assumptions are quite simple yet informative. The shocks to permanent income identified by the model explain most of the forecast-error variance in consumption, even at the one-quarter-ahead forecast horizon. Temporary factors explain most of the near- and medium-term forecast-error variance in spending on durable goods and total income. Nonetheless, the permanent shocks account for between one-fifth and one-third of the variation in these variables at the two-year forecast horizon, and the majority of the variance beyond 20 years.

Small-scale econometric models produce easily interpretable results that are robust to a number of specification issues which can plague large-scale systems. However, small models are too simple to address many of the issues faced by researchers and policymakers.

FIGURE 8

The output gap and its forecasted values as of 2005:Q2

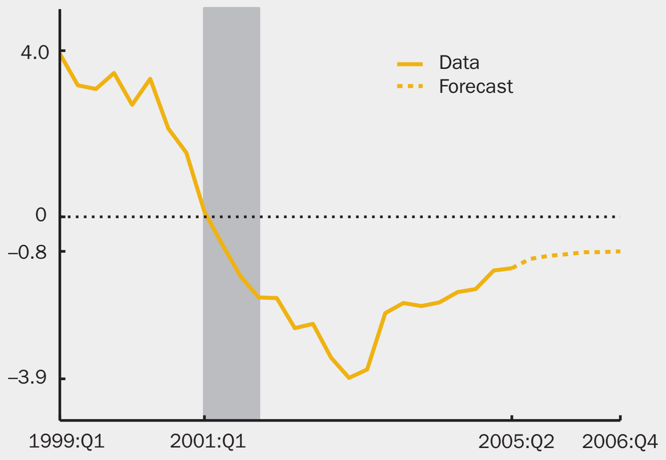

Notes: The shaded area is a recession as identified by the National Bureau of Economic Research. See the text for further details.

Accordingly, our future research will explore extensions to our model that can address interesting economic questions without adding overly restrictive identifying assumptions or burdensome model complexity.

One extension is to include the number of hours worked in the economy. Policymakers often are concerned about the implications of labor market slack for inflation, and transitory movements in hours are one proxy for such slack. Real business cycle researchers argue that changes in households' allocation of time between work and leisure explains much of the variation in economic output. Thus, movements in hours may help forecast both inflation and output.

Second, we plan to add long-term interest rates to our analysis. Many physical assets are illiquid, and households are averse to taking on risky investment ventures. As a result, investment projects may be financed with long-term borrowing, and the interest rates on these loans will not be simple averages of current and future one-period rates. Furthermore, inflation is a key determinant of the real value of money, particularly of funds that are committed for an extended period of time. Accordingly, long-term interest rates may be a useful predictor of both real investment and inflation.

Third, some technological change is embodied in new capital machinery. For example, advances in computing power are embodied in new computer chips and improvements in fuel efficiency are embodied in new jet engines. This type of technological change will be reflected in a decline in the price of capital goods relative to the price of consumption. We plan to follow Fisher (2005) and include these relative prices, thereby allowing the model to distinguish between embodied and disembodied sources of technological progress. We speculate that this distinction will result in more informative estimates of the output gap.

Finally, following Taylor (1993), a large literature relates changes in the federal funds rate to the output gap and differences between actual inflation and the Federal Reserve's target. We plan to investigate if the Taylor rule based on the output gap generated by our model provides any different interpretation of policy than relationships based on other measures of the output gap.

NOTES

[1]Our model is in the spirit of Cochrane (1994). He estimates a vector autoregression in the growth rates of consumption and gross national product (GNP). The model also includes the lagged log ratio of consumption to GNP as an explanatory variable. Because the growth rates are stationary, the inclusion of the lagged log ratio forces the impact of the model's shocks on the levels of consumption and GNP to be the same in the long run. Cochrane then identifies permanent and transitory shocks based on the restriction that shocks that do not contemporaneously affect consumption must be transitory.

[2]The distinction between durable goods and the service flow from those goods is subtle, and requires that we make some adjustments to the data published in the *National Income and Product Accounts* from the U.S. Bureau of Economic Analysis. We return to this issue in the data subsection.

[3]This equals the average of many independent realizations of Z_t. Of course, there is only one actual realization of any given variable Z_t. Many independent realizations can only be generated hypothetically.

[4]Using the unified budget constraint in equation 2 to replace the sequence of one-period budget constraints in equation 1 requires ruling out the possibility that the household finances its expenditures with a Ponzi scheme.

[5]Any sequence of nondurable consumption growing at this rate paired with the assumed interest rate satisfies equation 3, the condition for the optimal allocation of consumption between today and tomorrow. This growth rate, C_0, and the condition for optimal durable goods consumption determines a path for S_t, which grows at the same rate. With these paths in place, the budget constraints in equation 1 determine the path for B_t.

[6]Recall that Y_t grows at the rate μ and that the one-period interest rate equals $\beta^{-1}\mu$, so the present value of the household's labor income equals $\sum_{t=0}^{\infty}(1/R)^t Y_0\mu^t = \sum_{t=0}^{\infty}\beta^t Y_0 = Y_0/(1-\beta)$.

[7]We construct our measure of real consumption of nondurables and services as the chain aggregate of NIPA consumption expenditures on nondurables, services, and the negative of expenditures on housing services.

[8]The realized errors from forecasting nondurable consumption must have permanent effects on its level in order for C_t to provide useful information about the economic response to changes in permanent income. In statistical terms, this means that C_t must have a stochastic trend. Formal statistical tests (not shown) provide no evidence against this assumption.

[9]Stock and Watson (2002) review this literature.

[10]We later discuss how to use forecasts of these deviations to recover forecasts of the nominal ratios' levels.

[11]Note that this estimator is an instrumental variables technique; the conditions in equation 9 are identical to necessary conditions identifying the lagged variables in the VAR as valid instruments.

[12]This forecast is constructed by taking the estimated coefficients from the model, the lagged data for 1983:Q4 and earlier, and simulating the model forward with all of the error terms set equal to zero. The VAR's growth rate and log ratio forecasts are then transformed to the log levels shown in the figure. Box 2 (p. 66) gives more details on generating forecasts from our model.

[13]Cochrane (1994) finds his transitory shock is more important than the permanent shock for forecasting gross national product for about a year. After that, the permanent shock is more important.

[14]See Kuttner (1994) for one solution to this problem based on an unobserved components model.

REFERENCES

Campbell, Jeffrey R., and Zvi Hercowitz, 2004, "The role of collateralized household debt in macroeconomic stabilization," Federal Reserve Bank of Chicago, working paper, No. WP-2004-24.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, 1999, "Monetary policy shocks: What have we learned and to what end?," in *Handbook of Macroeconomics*, Vol. 1A, John B. Taylor and Michael Woodford (eds.), Amsterdam: Elsevier Science B.V., chapter 2.

Cochrane, John H., 1994, "Permanent and transitory components of GNP and stock prices," *Quarterly Journal of Economics*, Vol. 109, No. 1, pp. 241–265.

Fisher, Jonas D.M., 2005, "The dynamic effects of neutral and investment-specific technology shocks," Federal Reserve Bank of Chicago, manuscript.

Friedman, Milton A., 1957, *A Theory of the Consumption Function*, Princeton, NJ: Princeton University Press.

Hall, Robert E., 2005, "Separating the business cycle from other economic fluctuations," in *The Greenspan Era: Lessons for the Future*, Federal Reserve Bank of Kansas City, forthcoming.

__________, 1978, "Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence," *Journal of Political Economy*, Vol. 86, No. 6, pp. 971–987.

Hansen, Lars P., 1982, "Large sample properties of generalized method of moments estimators," *Econometrica*, Vol. 50, No. 4, pp. 1029–1052.

Kahn, James A., Margaret M. McConnell, and Gabriel Perez-Quiros, 2002, "On the causes of the increased stability of the U.S. economy," *Economic Policy Review*, Federal Reserve Bank of New York, Vol. 8, No. 1, May, pp. 183–202.

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo, 1988, "Production, growth and business cycles I: The basic neoclassical model," *Journal of Monetary Economics*, Vol. 21, No. 2–3, March/May, pp. 195–232.

Kuttner, Kenneth N., 1994, "Estimating potential output as a latent variable," *Journal of Business and Economic Statistics*, Vol. 12, No. 3, pp. 361–368.

McConnell, Margaret M., and Gabriel Perez-Quiros, 2000, "Output fluctuations in the United States: What has changed since the early 1980s?" *American Economic Review*, Vol. 90, No. 5, pp. 1464–1476.

Sims, Christopher A., 1980, "Macroeconomics and reality," *Econometrica*, Vol. 48, No. 1, pp. 1–48.

Stock, James H., and Mark W. Watson, 2002, "Has the business cycle changed and why?," in *NBER Macroeconomics Annual 2002*, Mark Gertler and Kenneth Rogoff (eds.), Cambridge, MA: National Bureau of Economic Research, pp. 159–218.

Taylor, John B., 1993, "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy,* Vol. 39, pp. 195–214.