Consumption, Savings, and the Meaning of the Wealth Effect in General Equilibrium

Carl D. Lantz and Pierre-Daniel G. Sarte

Over the latter half of the 1990s, the U.S. economy experienced both a substantial decrease in the savings rate and a significant run-up in household net worth. Between 1994 and 2000, the gross private savings rate fell from 17 to 12 percent, while the personal savings rate declined from above 6 percent to less than zero. Over the same period, the value of household sector equity holdings (including those owned by nonprofits, pensions, and other fiduciaries) increased nearly 150 percent for a dollar gain in excess of $6 trillion.

At some level, the decline in savings and the rise in household equity value during that period appeared to point towards a strengthening of the economy. According to the Permanent Income Hypothesis (PIH), households save less in a given period if they expect future increases in their income. Along these lines, the dramatic gain in stock market wealth was thought to partly reflect future opportunities made available to firms by rapid advances in information technology. Both the fall in savings and the rise in net wealth seemed consistent with the rapid growth of consumption during that period.

Despite the rosy outlook implied by the PIH at the close of the decade, the U.S. economy slowed down considerably in 2000. Specifically, the growth rate of per capita consumption fell to 2 percent in the first quarter of 2001 from nearly 7 percent in the same quarter of the previous year. Between the first quarter of 2000 and that of 2001, household net worth fell by 8 percent, or

The views expressed in this article are the authors' and do not necessarily represent those of the Federal Reserve Bank of Richmond or the Federal Reserve System. We wish to thank Michael Dotsey, Margarida Duarte, Thomas Humphrey, and Yash Mehra for helpful suggestions. All errors are our own.

\$3.5 trillion. In light of these developments, it seems only natural to question the significance of the data in the late 1990s. With this question in mind, this article seeks to emphasize the following points.

First, the PIH notwithstanding, a fall in savings today may not necessarily reflect expected future gains in income, but rather the current realization of a negative economic shock. Within the context of a simple neoclassical growth model with investment adjustment costs, we show that an unanticipated permanent fall in productivity leads to a contemporaneous fall in both consumption and savings. The fall in savings continues several periods into the future and a lower steady-state level of savings ultimately emerges. It remains true, in this model, that a fully anticipated increase in future productivity also leads to a contemporaneous fall in savings as households seek to smooth consumption. In the latter case, however, the savings rate eventually reaches a higher steady state level as the shock is realized.

Second, it is important to recognize that discussions of the wealth effect, such as those in Ludvigson and Steindel (1999) or Mehra (2001), are often carried out in a partial equilibrium setting. In such a setting, both the rate of interest and the level of wealth are exogenous with respect to contemporaneous consumption (i.e., wealth is a state variable). In contrast, general equilibrium considerations imply that wealth, the rate of interest, and consumption all contemporaneously react to the various disturbances affecting the economy. Thus, an unanticipated permanent increase in productivity leads to a simultaneous rise in both consumption and household net worth. Note, however, that consumption does not respond directly to wealth. Rather, both variables react simultaneously to the higher level of productivity. The implication of this dual reaction is that the *measured* marginal propensity to consume out of wealth is unlikely to be constant, as is often assumed. Indeed, empirical studies such as those in Mehra (2001) and Ludvigson and Steindel (1999) have found that the magnitude of the wealth effect is dependent on the sample period in question. This lack of time consistency in the wealth parameter would be expected if the nature of the shocks impacting the economy was changing over different sample periods.

In general, it can be misleading to think in terms of households' marginal propensity to consume out of wealth. Such thinking presumes that important movements in wealth exist that are independent of economic fundamentals. However, the value of corporate equity reflects the present discounted value of future firm dividends and, in a general equilibrium framework, both the discount rate and dividends respond to changes in the economic environment.

To make matters concrete, we show that consumption and wealth can move in opposite directions in some cases. When a future increase in productivity is fully anticipated, at the time of anticipation consumption rises while the value of household equity falls. Although households eventually hold more wealth in the new steady state, the initial fall in equity value reflects higher future discount rates consistent with the anticipated increase in productivity. A partial equilibrium framework prohibits this finding from ever arising because the rate of interest is held fixed.[1]

In this article, we first present some basic empirical facts regarding consumption, savings, and wealth in U.S. data. We next outline a simple theoretical framework that allows us to simultaneously explore the price of corporate equity and households' consumption-savings decisions. Finally, we analyze the results from several numerical experiments related to both anticipated and unanticipated shocks to total factor productivity.

1. CONSUMPTION AND THE SAVINGS RATE IN U.S. DATA

Figure 1 shows the behavior of two alternate measures of the U.S. savings rate over the past 41 years. Panel a of Figure 1 captures the most basic National Income and Product Accounts (NIPA) measure of savings, Personal Disposable Income less Personal Consumption Expenditures in 1996 dollars. The savings rate in panel b is computed using Gross Private Savings which, in addition to Personal Savings, includes retained earnings by firms. We can see that both measures of the savings rate fell drastically over the 1990s and, by early 2001, had reached their lowest recorded levels.

We suggested earlier that a desire to smooth consumption may lead households to save less today if they expect future gains in their income or, alternatively, to save more if they expect future declines in their income. In particular, Hall (1978) argued that the consumption behavior of a household at a given date was based on all of that household's future discounted earnings. Milton Friedman (1957) was perhaps the first to draw a distinction between changes in permanent and transitory income. Figure 2 illustrates (normalized) movements in the savings rate four quarters prior to each of the past five U.S. recessions. In panel a, we can see that the personal savings rate generally rises during the year prior to a recession. However, this tendency is not clear-cut. Moreover, it is much less pronounced for the gross private savings rate in panel b. In this case, in the four quarters preceding two of five recessions, the savings rate either falls or remains the same. Figure 3 plots the cross-correlations between our two measures of the savings rate and output at different leads and lags. Both the personal savings rate and the gross private savings rate show a negative correlation with future values of GDP. Hence, there seems to be some evidence to support the PIH. However, the magnitude of the cross-correlations

[1] See Kiley (2000) for a more detailed description of stock price behavior in a production economy versus a partial equilibrium setting.

Figure 1 Measures of Savings

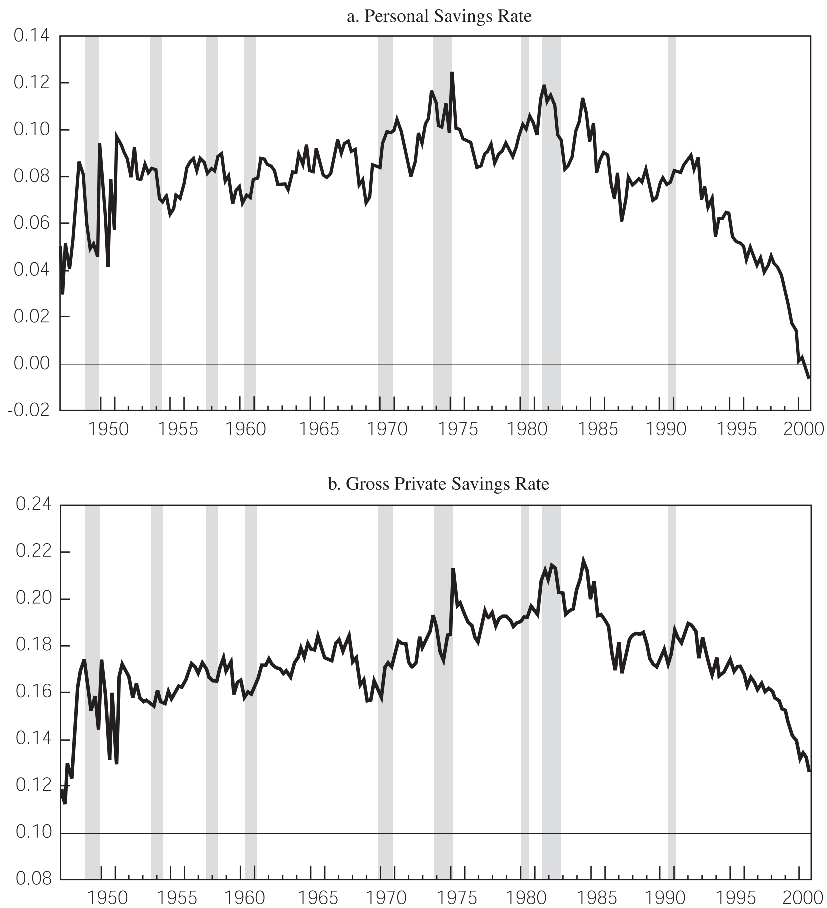

shown in Figure 3 is relatively low, and it is possible that factors other than expectations of future changes in income help drive the behavior of the savings rate.

2. A SIMPLE THEORETICAL PERSPECTIVE

In order to explore some of the issues introduced above, we now describe a model that can be simultaneously used to price corporate equity and address household consumption-savings decisions. For simplicity, we abstract from the inclusion of a noncorporate sector and intangible assets, as well as several aspects of the U.S. tax system. McGrattan and Prescott (2000), however, suggest that these considerations are important in calibration exercises meant

Figure 2 Savings Rate and Equity Price Behavior Prior to Various U.S. Recessions

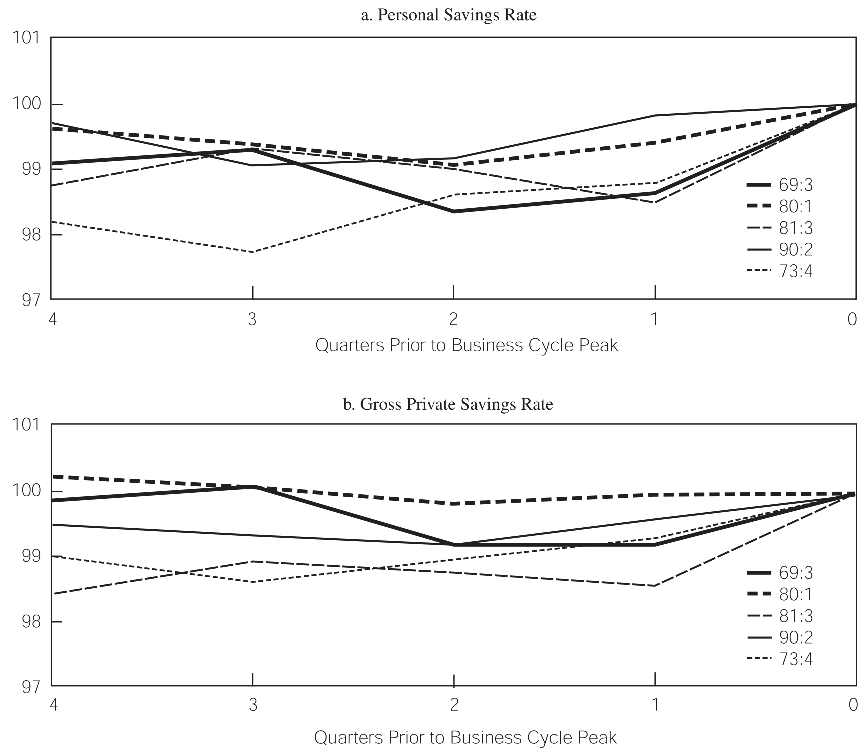

to match data from the NIPA and the *Statistics of Income* (SOI). In particular, the authors argue that the historical behavior of asset prices and returns can be largely explained by changes in tax and regulatory policies as well as by the evolution of the institutions affecting asset markets.

In this model, the economic environment consists of a large number of identical households and firms. Each firm has access to a constant returns technology,

$$y_t = z_t k_t^{\alpha} n_t^{1-\alpha}, \; 0 < \alpha < 1, \tag{1}$$

where y_t is the firm's output at a given date t, n_t denotes labor input, z_t is a random technological shift parameter, and k_t represents the firm's capital stock. In this article, we shall think of firms as owning their capital stock instead of renting it from households. Households will be thought of as owning claims on firms' net cash flows, e.g., equity shares.

Figure 3 HP Filtered Cross-Correlations with Output corr$[x_t, y_{t+k}]$

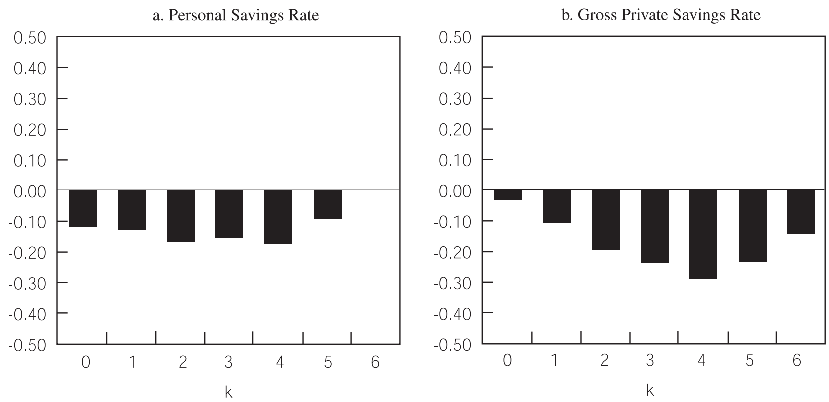

Barro and Sala-i-Martin (1995) suggest that if the stock of capital includes a human component, then one will anticipate substantial adjustment costs in investment. According to the authors, "the learning process takes time, and attempts to accelerate the training process are likely to encounter rapidly diminishing rates of return" (p. 119). Hence, we model the evolution of a firm's capital stock as

$$k_{t+1} = (1 - \delta)k_t + \phi\left(\frac{i_t}{k_t}\right)k_t, \tag{2}$$

where $0 < \delta < 1$ is the capital depreciation rate and i_t represents the firm's investment decision at date t. The function $\phi(\cdot)$, with $\phi'(\cdot) > 0$, captures the idea of adjustment costs in investment. Thus, the higher the level of investment relative to the current capital stock, the more costly it becomes to increase next period's capital. Observe that the function $\phi''(\cdot) < 0$ indexes the degree to which adding to the capital stock becomes costly.[2] In addition, note also that the book value of capital at date t, k_t, reflects investment decisions made at date $t-1$. Therefore, k_t cannot respond contemporaneously to changes in the economic environment. In contrast, if we think of the firm as having a fixed number of equity shares outstanding, the value of these shares can contemporaneously react to disturbances affecting the economy. Put another way, we expect both household net worth and consumption to move simultaneously in response to various shocks.

[2] For an early discussion of this formulation of investment adjustment costs, see Abel and Blanchard (1983).

Firms pay each unit of labor the wage rate w_t, and their net cash flow at t is consequently given by

$$z_t k_t^{\alpha} n_t^{1-\alpha} - i_t - w_t n_t. \tag{3}$$

We assume that this cash flow is paid to households in the form of dividends, D_t. Each firm attempts to maximize the present discounted value of future profits. The representative firm's problem, therefore, can be summarized as

$$\max \sum_{\tau=0}^{\infty} \Pi_{i=-1}^{\tau-1} Q_{t+i} \left[z_{t+\tau} k_{t+\tau}^{\alpha} n_{t+\tau}^{1-\alpha} - i_{t+\tau} - w_{t+\tau} n_{t+\tau} \right], \tag{P1}$$

subject to the sequence of constraints given by (2). In (P1), Q_{t-1} denotes the price of a security that pays one unit of the consumption good at date t.

The solution to the firm's problem must satisfy the following first-order conditions,

$$w_t = (1-\alpha) z_t k_t^{\alpha} n_t^{-\alpha}, \tag{4}$$

$$\lambda_t \phi' \left(\frac{i_t}{k_t} \right) = 1, \tag{5}$$

and

$$\begin{aligned} & Q_t \alpha z_{t+1} k_{t+1}^{\alpha-1} n_{t+1}^{1-\alpha} \\ = \; & \lambda_t - Q_t \lambda_{t+1} \left[(1-\delta) + \phi \left(\frac{i_{t+1}}{k_{t+1}} \right) - \phi' \left(\frac{i_{t+1}}{k_{t+1}} \right) \frac{i_{t+1}}{k_{t+1}} \right], \end{aligned} \tag{6}$$

where $\lambda_t \geq 0$ is the Lagrange multiplier associated with (2). Equation (4) simply equates the wage rate to the marginal product of labor. Equation (5) suggests that it is optimal for the firm to invest up to the point where the cost of one additional unit of investment (in terms of foregone profits) exactly offsets the marginal gain from increasing next period's capital stock.

As mentioned earlier, the representative household owns all firms and receives their profits, D_t, as dividends. At date t, the typical household's net worth, A_t, consists of stock market wealth and bonds. Specifically, we denote the market value of household equity by $V_t X_t$, where V_t represents the price of firms' outstanding equity shares and X_t is the number of shares held by the household. Agents also own one-period bonds, B_t, where a bond purchased at date t pays one unit of the consumption good at time $t+1$. The representative household maximizes its lifetime utility and solves

$$\max \sum_{\tau=0}^{\infty} \beta^{\tau} \frac{c_{t+\tau}^{1-\sigma} - 1}{1-\sigma}, \; \sigma > 0, \tag{P2}$$

subject to the sequence of constraints

$$c_t + V_t X_{t+1} + Q_t B_{t+1} = (V_t + D_t) X_t + B_t + w_t n_t. \tag{7}$$

Household income on the right-hand side of equation (7) stems from the ownership of firms, with dividend earnings given by $D_t X_t$, earnings from bonds, B_t, and labor income, $w_t n_t$. These earnings can be used to purchase consumption goods, new equity shares, and bonds. The first-order conditions associated with the household problem are

$$c_t^{-\sigma} = \psi_t, \tag{8}$$

$$Q_t = \beta \left(\frac{\psi_{t+1}}{\psi_t} \right), \tag{9}$$

and

$$V_t = \beta \left\{ \left(\frac{\psi_{t+1}}{\psi_t} \right) \left[V_{t+1} + D_{t+1} \right] \right\}, \tag{10}$$

where ψ_t is the multiplier associated with the household budget constraint (7). Note that equations (9) and (10) can be used together to yield

$$V_t = \sum_{\tau=1}^{\infty} \Pi_{i=1}^{\tau-1} Q_{t+i} D_{t+\tau}. \tag{11}$$

In other words, the price of a firm's outstanding equity shares reflects the expected present discounted value of its future dividends. In this model, therefore, even shocks that affect only future profit opportunities and discount rates will lead to changes in today's household wealth.

Observe that the multiplier λ_t in (5) can be interpreted as the shadow price of installed capital. In particular, the Appendix shows that equations (6) and (11) can be used to derive

$$V_t = \lambda_t k_{t+1}. \tag{12}$$

Since $\phi''(.) < 0$, an increase in investment leads to a rise in λ_t by equation (5), as well as a rise in k_{t+1}. Hence, in thinking about the effects of various shocks below, we need only keep track of the investment response in order to understand movements in the value of corporate equity.[3]

An equilibrium for the economy we have just presented must satisfy firms' optimality conditions (4) through (6), as well as households' optimality conditions (8) through (10). In addition, the goods market clearing condition,

$$c_t + i_t = y_t, \tag{13}$$

must hold. In equilibrium, we further have that $X_t = X_{t-1} = 1$ for all t and, since households are identical, bonds are in zero net supply, $B_t = 0$ for all t. Equation (13) implies that savings equals investment, $s_t = y_t - c_t = i_t$.

[3] Hayashi (1982) shows that equation (12) always holds when the production technology is constant returns to scale.

Before investigating the joint response of consumption, savings, and wealth to different changes in the economic environment, we must first assign values to the exogenous parameters of our model. Each period represents a quarter, and we set δ and σ to 0.025 and 2 respectively. These values for δ and σ are standard in quantitative studies of business cycles. In the steady state, equations (9) and (11) imply that the price-earnings ratio, V/D, is given by $\beta/(1-\beta)$. Hence, we set β to 0.983 in order to generate a long-run annualized price-earnings ratio of 14.5.[4] We set α to 1/3 which leads to an investment share in output of 20 percent in the steady state. Finally, we set the parameter that governs the degree of adjustment costs, ϕ'', to -10. This calibration implies that the elasticity of the investment:capital ratio with respect to Tobin's q is approximately 5. Baxter and Crucini (1993) explore a variety of possible calibrations for this elasticity parameter, ranging from 1 to 15, without substantially altering their results. This remains true in our framework.

On the Significance of the Wealth Effect in General Equilibrium

The solution to the model above implies a law of motion for the vector of state variables, s_{t+1} as a function of s_t, where s_t consists of the capital stock, k_t, and the random technological shift parameter, z_t. This solution also links control variables, such as consumption, c_t, and the market capitalization of firms, V_t, to the state variables. Therefore, in a linearized form, we have

$$c_t = c_0 + c_k k_t + c_z z_t \tag{14}$$

and

$$V_t = v_0 + v_k k_t + v_z z_t, \tag{15}$$

where c_0, $v_0, \ldots$ are functions of the deep parameters of the model capturing preferences and technology. Solving for k_t in equation (15) and substituting the resulting expression in (14) yields

$$c_t = \underbrace{(c_0 - \frac{c_k}{v_k} v_0)}_{\text{constant}} + \underbrace{(\frac{c_k}{v_k})}_{\beta} V_t + \underbrace{(c_z - \frac{c_k}{v_k} v_z) z_t}_{u_t}. \tag{16}$$

This last equation often forms the basis of regression equations that are meant to uncover the size of the wealth effect, $\partial c_t/\partial V_t = c_k/v_k = \beta$. Observe that the only source of random disturbances in equation (16) stems from movements in productivity, z_t. Moreover, because changes in equity V_t are necessarily correlated with changes in fundamentals, z_t, it will be important to make

[4] Until recently, this value has been approximately the average implied by the S&P 500 index since 1949.

use of instrumental variables to properly estimate the coefficient β. That being said, since all movements in both c_t and V_t are generated from changes in economic fundamentals, estimates of the marginal propensity to consume out of wealth are of little use in this environment. More to the point, the expression $\partial c_t/\partial V_t$ is meaningful only to the degree that there exist significant exogenous movements in net worth, ∂V_t, that are unrelated to changes in underlying economic conditions. Such movements may reflect, for example, the existence of stock market bubbles. In our environment, however, changes in consumption and wealth are necessarily linked through movements in productivity and given by

$$\begin{aligned} \frac{\partial c_t}{\partial z_t} &= \left(\frac{c_k}{c_v}\right)\frac{\partial V_t}{\partial z_t} + \left(c_z - \frac{c_k}{v_k}v_z\right) \\ &= c_z. \end{aligned} \tag{17}$$

3. NUMERICAL EXAMPLES

We will now explore the behavior of our economy when the underlying source of uncertainty lies in total factor productivity, z_t. We shall examine the effects of both unanticipated and anticipated changes in productivity, and outline significant differences in the way the economy reacts to these shocks. To emphasize these differences, we shall also compute the cross-correlations of consumption and the savings rate with stock market wealth under both these parameterizations of productivity shocks.

The Effects of Unanticipated Shocks in Productivity

Figure 4, panel a, depicts an unanticipated and permanent 1 percent fall in productivity. As a result of this shock, output falls immediately as depicted in Figure 4, panel d, and continues falling towards a lower steady state value. Observe that both consumption and savings mimic the output response. Both variables fall at the time of the shock and eventually reach a lower steady state level. In this case, therefore, a fall in savings does not indicate better times ahead, as a naive interpretation of the PIH suggests. Instead, by allowing households to consume some of their capital, diminished savings behavior softens the fall in consumption. It remains true, of course, that the economy is unambiguously worse off in the long run.

In this numerical experiment, the savings rate decreases dramatically on impact and then *rises* on its way to the final steady state. This is shown in Figure 5, panel b. In the new long-run equilibrium, however, the savings rate is ultimately lower relative to its level in the period prior to the shock. This example suggests that it may be difficult to identify the source of a given decline in the savings rate in the data. In particular, we shall see below that one version of the PIH continues to hold in general equilibrium. That is,

Figure 4

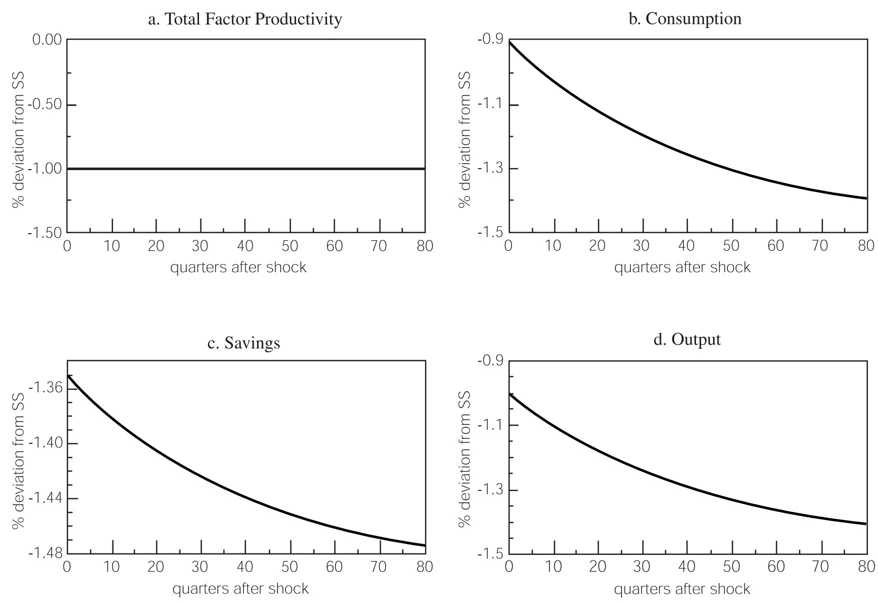

an anticipated increase in future productivity also leads to a decrease in the savings rate today, followed by a gradually increasing path. In the case of this anticipated increase, however, the savings rate eventually increases all the way to a *higher* steady state level.

Figure 5 also shows that the interest rate, firms' dividends, and the market value of equity all decrease when the negative productivity shock is realized. Given equation (12), the fall in equity is relatively easy to follow. Because the level of savings falls in response to the shock, firms are forced to cut back on investment, which directly leads to a decrease in the value of corporate equity. Note that this decline in equity is consistent with the fall in aggregate dividends in Figure 5, panel c, but is mitigated by the decrease in interest rates during the transition to the new steady state. Since the rate of interest is simply the inverse of Q_t in equation (9), the steady fall in consumption in Figure 4, panel b, indeed implies a decline in interest rates until the new long-run equilibrium is reached.

Finally, in this example, Figures 4b and 5d show that consumption and wealth respond to the shock in the same direction. As we have already pointed out, however, it should be clear that there is no sense in which consumption responds directly to movements in wealth. Furthermore, the nonlinearity of the impulse responses implies that the measured marginal propensity to consume out of wealth will not be constant in this case. This implication is at variance

Figure 5

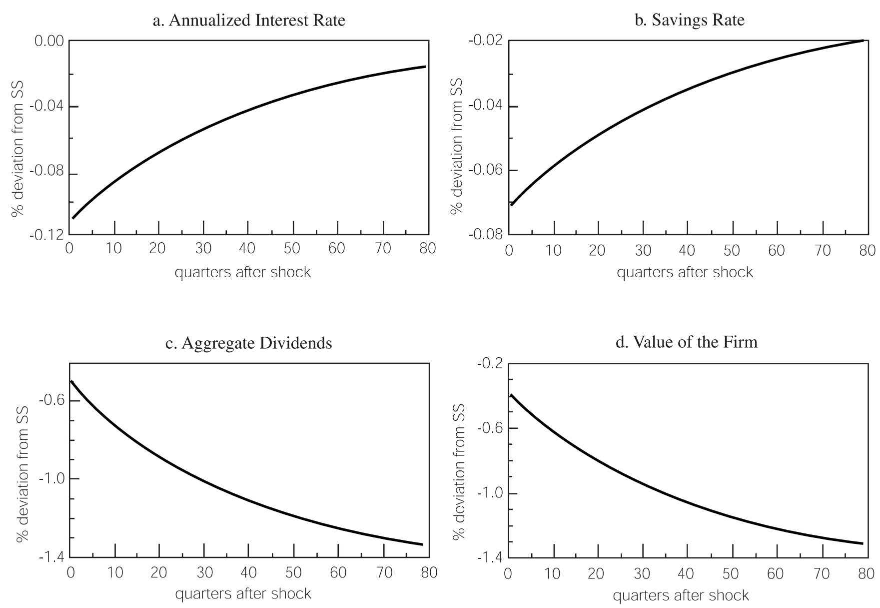

with studies, such as Davis and Palumbo (2001) and Poterba and Samwick (1995), that have attempted to measure the additional increase in consumption stemming from a rise in household equity.

The Effects of Anticipated Changes in Productivity

We now study the model economy's response to an anticipated permanent positive shock to total factor productivity. One interpretation of such a shock may involve the conception of a new technology whose actual implementation is likely to take time. We shall see that in the short run, there exist similarities in the way savings respond to an anticipated positive shock and an unanticipated negative shock. These similarities, while they can make the interpretation of savings data ambiguous at times, eventually dissipate in the long run.

Figure 6, panel a, depicts a 1 percent positive shock in total factor productivity that takes place four periods in the future. This shock, however, is fully anticipated by both households and firms in the current period. Because productivity, and thus output, is expected to increase, household consumption immediately rises in Figure 6, panel b. This response reflects a desire to smooth consumption that is implicit in the household problem. However, since the capital stock, k_t, is fixed at time zero, output cannot change at the

Figure 6

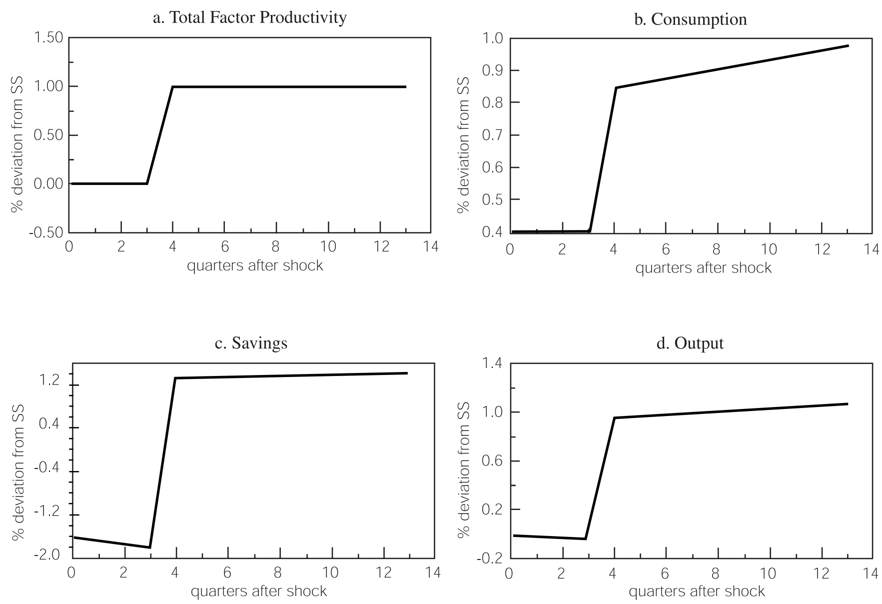

time of the shock. It must be the case, therefore, that savings initially fall in a way consistent with the PIH, as shown in Figure 6c.

Observe that because the initial increase in consumption is sustained until the productivity shock takes place, the level of savings continues to fall in the short run. Therefore, as households find it optimal to temporarily consume part of the capital stock, output declines between period 0 and period 4. Once the positive productivity shock occurs in period 4, consumption, savings, and output all increase and begin converging towards their new steady state. In our context, adjustment costs limit the extent to which households wish to increase consumption initially. To be specific, since firms will find it optimal to increase investment once the shock occurs, and the marginal product of capital will consequently rise, it will be important that the capital stock not be too low at the point of the shock. Recall that the nature of investment adjustment costs is such that the higher the level of investment relative to the current capital stock, the more costly it becomes to increase the next period's capital.

Figure 7

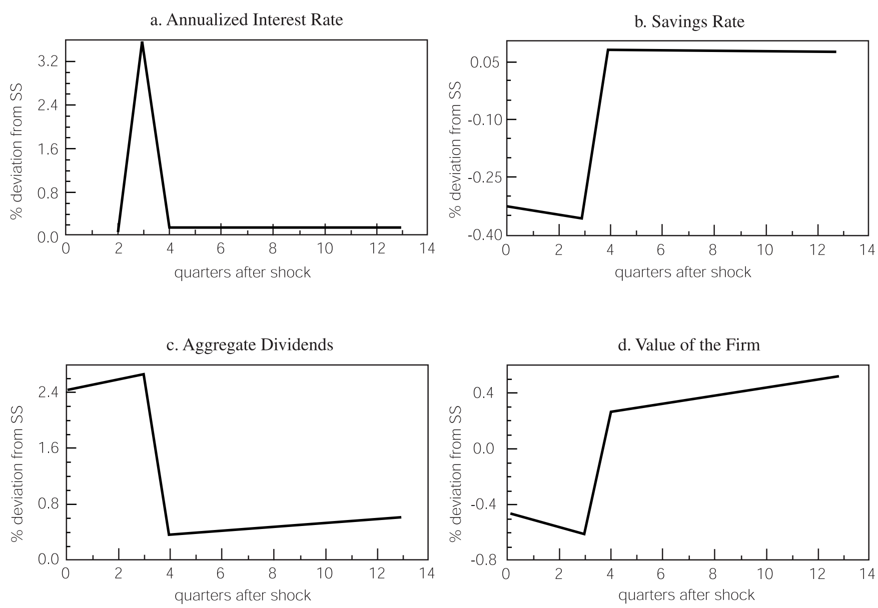

Figure 7, panel d, shows that the value of corporate equity actually falls when the productivity shock is anticipated at time zero. This result can be most easily understood in terms of the fall in savings in Figure 6c and the resulting decline in investment. More importantly, this finding clearly indicates that consumption, as shown in Figure 6b, and wealth do not have to move in the same direction. This result is at odds with many empirical studies in which consumption always responds positively to wealth within the assumed theoretical framework. On a related note, the impulse responses depicted in Figures 6 and 7 suggest that the data in the late 1990s were not necessarily indicative of a future strengthening of the economy. As we pointed out in our introduction, both consumption and wealth rose during that period while savings fell. Our numerical experiment suggests that an anticipated positive shock to productivity, while leading to a fall in savings and a rise in consumption during the current period, generates a fall in wealth initially.

Finally, Figure 7, panel a, illustrates a remarkable increase in the interest rate in the period prior to the realization of the shock. This noticeable increase is consistent with the jump in consumption that occurs in the next period when the shock takes place. In particular, the high rate of interest prevents consumption from rising too dramatically in anticipation of the productivity increase. Moreover, observe that the interest rate spike is also consistent with

Figure 8 Model-Generated Cross-Correlations with Wealth
corr$[x_t, V_{t+k}]$

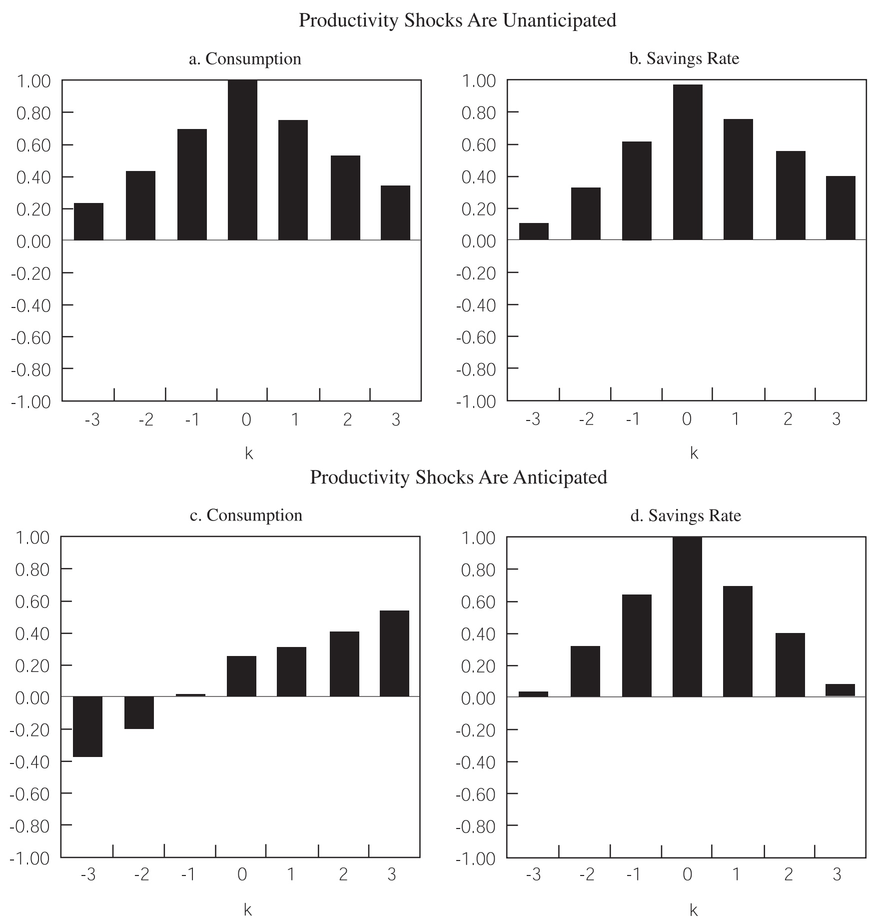

the initial fall in wealth in Figure 7c. Once the shock has occurred, the high rate of interest depicted in Figure 7a is no longer part of the present discounted value calculation with respect to future earnings. As a result, the value of corporate equity increases markedly.

Implied Cross-Correlations between Consumption, Savings, and Wealth

Thus far, we have seen that the nature of productivity shocks, whether they are anticipated or unanticipated, has significant implications for the reactions of key economic variables. In particular, we have seen that wealth and consumption do not always have to respond in the same direction to a given productivity shock. We will emphasize this point below by showing important differences in the cross-correlation pattern of the data generated under each type of shock.

Figure 8 presents the cross-correlations of consumption and the savings rate with stock market wealth generated by the model. As in the real-business-cycle literature, we first assume (in Figures 8a and 8b) that the dominant source of uncertainty lies in productivity shocks, which we calibrate as

$$\ln z_t = \rho_z \ln z_{t-1} + \varepsilon_{zt}, \tag{18}$$

where $\rho_z = 0.95$ and ε_{zt} is an i.i.d. normal random variable with mean zero and standard deviation 0.01. The model statistics depicted in Figure 8 are the mean values calculated from 200 simulations of samples with 216 observations each, the number of quarterly observations in postwar U.S. data. Figures 8c and 8d present the same cross-correlations under the assumption that all productivity shocks are anticipated four periods in advance.

As we can see from the simulations in Figure 8, the cross-correlation patterns of consumption and savings with wealth are quite different depending on the nature of productivity shocks. When shocks are unanticipated, the contemporaneous correlation between consumption and wealth is very near 1. This contemporaneous correlation, however, is much lower at 0.25 when productivity shocks are anticipated. Therefore, to the degree that the U.S. economy is continuously hit by a variety of shocks that are both unanticipated and anticipated—to technology, preferences, or even public expenditures—and whose processes may have changed over time, it is unlikely that a regression of consumption on wealth would uncover a stable coefficient over different sample periods.

Finally, it is important to recognize that the cross-correlation patterns depicted in Figure 8 may change significantly with the particular model at hand. For instance, Constantinides (1990) and Abel (1990) suggest that habit formation is an important factor in explaining consumption behavior. When subject to habit formation, consumption reacts to various shocks only with a lag, and this lag may be essential in helping us understand U.S. consumption data. In addition, the model we have examined does not allow for the presence of credit-constrained households. For these households, consumption may be more tied to current income and wealth than is suggested by permanent income households.

4. CONCLUDING REMARKS

At the close of the 1990s, the U.S. economy experienced declining savings, a rise in household equity value, and rapidly growing consumption. At some level, this data appeared indicative of a strengthening economy going forward. The Permanent Income Hypothesis (PIH) indeed suggests that savings should fall in the current period if increases in income are expected in the future and that the fall in savings would simply reflect households' desire to smooth consumption.

Contrary to this optimistic scenario, the U.S. economy slowed down considerably in 2000. Consequently, it seems natural to reevaluate the significance of the data in the late 1990s. With this task in mind, we have stressed the following points.

First, the PIH notwithstanding, a fall in savings does not necessarily reflect the expectation of future gains in income but can instead reflect the current realization of an unanticipated, negative economic shock. In the case of an unanticipated decline in productivity, the level of savings continues to fall until it reaches a lower steady state level. In contrast, in response to an anticipated positive shock to future productivity, savings eventually rise to a higher steady state level even if they fall initially.

Second, we have attempted to make clear that consumption and wealth simultaneously react to fundamental changes in the economic environment. In a general equilibrium context, there is no sense in which consumption responds directly and positively to changes in wealth. The latter notion has, in fact, been the starting point for many empirical studies, but we have shown that when a future increase in productivity is fully anticipated, consumption and wealth may initially move in *opposite* directions. Furthermore, because both the consumption and wealth responses to productivity disturbances are nonlinear, the measured marginal propensity to consume out of wealth is unlikely to be constant. In light of these results, the data on consumption, savings, and wealth in the late 1990s should not necessarily have been interpreted as presaging a future strengthening of the economy. Our numerical experiments suggest that an anticipated rise in productivity, while leading to a fall in savings and an increase in consumption in the current period, initially generates a short-run decline in wealth. The last response is at odds with the behavior of wealth at the end of the last decade.

APPENDIX: DERIVATION OF TOBIN'S q

This appendix describes the derivation of equation (12) in the text. Specifically, multiplying both sides of equation (6) by $k_{t+1} \geq 0$ yields

$$\begin{aligned} \lambda_t k_{t+1} &= Q_t \alpha y_{t+1} + Q_t \lambda_{t+1} \left[(1-\delta) k_{t+1} + \phi \left(\frac{i_{t+1}}{k_{t+1}} \right) k_{t+1} \right] \\ &\quad - Q_t \lambda_{t+1} \phi' \left(\frac{i_{t+1}}{k_{t+1}} \right) i_{t+1}. \end{aligned}$$

In this last expression, $\left[(1-\delta)k_{t+1} + \phi\,(i_{t+1}/k_{t+1})\,k_{t+1}\right]$ is simply k_{t+2} while $\lambda_{t+1}\phi'\,(i_{t+1}/k_{t+1}) = 1$ by equation (5). Therefore,

$$\lambda_t k_{t+1} = Q_t \underbrace{\left[y_{t+1} - w_{t+1} n_{t+1} - i_{t+1} \right]}_{D_{t+1}} + Q_t (s_{t+1}) \lambda_{t+1} k_{t+2}$$

since $\alpha y_{t+1} = y_{t+1} - w_{t+1} n_{t+1}$. By repeatedly substituting for $\lambda_{t+j}\, k_{t+j+1}$, $j \geq 1$, we have

$$\sum_{\tau=1}^{\infty} \Pi_{i=0}^{\tau-1} Q_{t+i} D_{t+\tau} = \lambda_t k_{t+1},$$

where $\sum_{\tau=1}^{\infty} \Pi_{i=0}^{\tau-1} Q_{t+i} D_{t+\tau}$ is simply V_t by equation (11) in the text. Thus, λ_t has the interpretation of Tobin's q.

REFERENCES

Abel, Andrew B. "Asset Prices under Habit Formation and Catching Up with the Joneses," *American Economic Review*, vol. 80 (May 1990), pp. 38–42.

______________, and Olivier J. Blanchard. "An Intertemporal Model of Saving and Investment," *Econometrica*, vol. 51 (May 1983), pp. 675–92.

Barro, Robert J., and Xavier Sala-i-Martin. *Economic Growth*. New York: McGraw-Hill, 1995.

Baxter, Marianne, and Mario J. Crucini. "Explaining Saving-Investment Correlations," *American Economic Review*, vol. 83 (June 1993), pp. 416–36.

Constantinides, George M. "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, vol. 98 (June 1990), 519–43.

Davis, Morris A., and Michael G. Palumbo. "A Primer on the Economics and Time Series Econometrics of Wealth Effects," Federal Reserve Board Finance and Economics Discussion Series Working Paper 2001-09 (January 2001).

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

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

Hayashi, Fumio. "The Permanent Income Hypothesis: Estimation and Testing by Instrumental Variables," *Journal of Political Economy*, vol. 90 (October 1982), pp. 895–916.

Kiley, Michael T. "Stock Price Fundamentals in a Production Economy," Federal Reserve Board Finance and Economics Discussion Series Working Paper (January 2000).

Ludvigson, Sydney, and Charles Steindel. "How Important Is the Stock Market Effect on Consumption?" Federal Reserve Bank of New York *Policy Review*, vol. 5 (July 1999), pp. 29–51.

McGrattan, Ellen R., and Edward C. Prescott. "Is the Stock Market Overvalued?" Federal Reserve Bank of Minneapolis *Quarterly Review*, vol. 24 (Fall 2000), pp. 20–40.

Mehra, Yash P. "The Wealth Effect in Empirical Life-Cycle Aggregate Consumption Equations," Federal Reserve Bank of Richmond *Economic Quarterly*, vol. 87 (Spring 2001), pp. 45–68.

Peek, Joe. "Capital Gains and Personal Savings Behavior," *Journal of Money, Credit and Banking*, vol. 15 (February 1983), pp. 1–23.

Poterba, James M. "Stock Market Wealth and Consumption," *Journal of Economic Perspectives*, vol. 14 (Spring 2000), pp. 99–118.

______________, and Andrew A. Samwick. "Stock Ownership Patterns, Stock Market Fluctuations, and Consumption," *Brookings Papers on Economic Activity*, 1995 (Issue 2), pp. 295–357.