Editors' Summary

THIS ISSUE of *Brookings Papers on Economic Activity* contains papers and discussions presented at the forty-fifth conference of the Brookings Panel on Economic Activity, which was held in Washington, D.C., on April 7 and 8, 1988. Two articles present new Keynesian theories of the macroeconomy and compare, against alternative theories, the ability of the new theories to explain key features of the economy. A third article examines the recent rapid growth of debt among U.S. nonfinancial corporations and its implications for their financial stability. A fourth article looks at the effect that capital market imperfections have on U.S. business investment and its financing. One report in this issue analyzes European unemployment and the likely consequences of demand stimulation in Europe. A second report looks at the extended slowdown in U.S. wage inflation.

IN TRADITIONAL Keynesian models, price and wage rigidities are the main reason why purely nominal shocks to aggregate demand result in real output and employment fluctuations. The absence of a rigorous explanation for such rigidities has made Keynesian models vulnerable to the charge of being ad hoc, and it has seemed anomalous that firms and workers would persist in sluggish wage and price behavior at the cost of substantial and sustained losses of employment and output. In recent years economists have taken up the challenge of providing a microeconomic explanation why nominal rigidities should exist, and how sluggish nominal adjustments might involve minor costs for a single firm, but major costs for the economy as a whole. In the first paper of this issue Laurence Ball, N. Gregory Mankiw, and David Romer summarize recent research that provides microeconomic foundations for nominal rigidities, and incorporate some of the central ideas of this research into a new Keynesian model of price adjustment. They analyze the relation between average inflation and the slope of the Phillips curve predicted by their

model, and provide an empirical test of that model against one based on the new classical tradition.

Imperfect competition plays an important role in new Keynesian models and eliminates several awkward implications of assuming perfect competition. Imperfect competition allows the firm effective control of its own price and, when combined with costs of adjustment, provides a microeconomic explanation why prices are rigid and output demand determined, and hence why deficient aggregate demand can cause unemployment. Imperfect competition explains why real wages need not be countercyclical, and, since equilibrium output is below the social optimum, also supports the common view that booms are welfare-improving.

These results follow from two theoretical observations about the nature of price rigidities based on "menu costs"—the costs of making price adjustments. The first theoretical observation is that, because firms set prices that are privately optimal when they make price changes, they suffer only second-order losses from not adjusting prices in response to small changes in nominal demand. Even if menu costs are small, it will not pay a firm to adjust its price to what would be the profit-maximizing level if those costs are ignored. As a result, changes in nominal demand, either up or down, have an effect on real output of the same order of magnitude as the change in demand. Because these changes occur in both directions, there is no presumptive change in the average output of the firm. Thus the only direct loss in either the firm's or society's welfare comes from the increase in the variability of the firm's output associated with price rigidities. These costs of variability in the firm's own output are of second order, both for the firm and for society.

The second theoretical observation is that there are aggregate demand externalities that make the variability cost greater for society than for the individual firm. Real income effects resulting from a firm's cutting or raising its price will be felt by workers and other firms in the economy. No one firm feels a noticeable part of this effect, so it does not influence its desire to adjust. But if all firms were to adjust their prices together, they would have a substantial effect on total aggregate demand. Because rigidity prevents this collective real income adjustment, the social costs of rigidity, even with symmetric demand shocks, may be much greater than the private costs.

Ball, Mankiw, and Romer observe that although these theoretical arguments explain how large and costly real fluctuations could in principle result from small frictions, the basic models from which these arguments are derived have trouble explaining the size and persistence of observed real output fluctuations. However, the authors report that including rigidities in real wages and prices together with asynchronized timing of price changes by different firms can amplify the effect on real output of rigidities in price and wage setting. Hence, they believe that such models can explain large and costly fluctuations in output.

In the second section of the paper the authors develop a specific, dynamic model of the economy with which to analyze the effect of menu costs. The model is similar in spirit to models by John Taylor and Olivier Blanchard describing the effect of wage rigidities, a major difference being that the timing of price changes in the authors' model is endogenous. The speed of adjustment of the aggregate price level depends on the frequency of price adjustment by individual firms, which is derived from profit maximization. As in the Taylor and Blanchard models, the firms' decisions are asynchronous, staggered evenly in time. A firm's profits are assumed to depend on its relative price, firm-specific shocks, and aggregate spending in the economy. The aggregate price level is simply the average of prices across firms. The crucial assumption is that there is a fixed cost of changing prices.

Although the formal analysis of the model is difficult—indeed, the authors are forced to find the equilibrium values of the model by numerical techniques—the qualitative results are intuitive and easily summarized. For a given frequency of price change, a nominal shock affects the price level gradually, passing fully into prices only asymptotically. The longer the interval between price changes, the slower the adjustment. As would be expected, the time interval itself decreases with higher inflation and with greater variances of aggregate and firm-specific shocks. As a consequence of this endogenous response of the interval between price shocks, the Phillips curve is steeper when any of these three variables is larger.

The authors show that the model generates quantitatively important effects when they specify reasonable values for its parameters. For example, an increase in the rate of inflation from zero to 10 percent, which decreases the average interval between price changes from 28

weeks to 19 weeks, steepens the short-run Phillips curve, reducing the fraction of a nominal shock that shows up as a real effect after six months from 0.5 to 0.3. An increase in the variability of aggregate demand from 3 percent to 10 percent, at an average inflation rate of 5 percent, has roughly the same effect on the steepness of the Phillips curve as raising the average rate of inflation from 5 percent to 15 percent.

This effect of variability in the authors' model is like that found in new classical models, but the reasons for it are different, and so also is the model's other main prediction. In Robert Lucas's new classical model, an increase in the variability of nominal aggregate demand causes firms to reduce their output response to nominal shocks because it reduces the informational value of prices to firms. But only variability matters in this model. The average rate of inflation is irrelevant to the output-inflation trade-off.

That there should be a systematic relationship between the average rate of inflation and the Phillips curve trade-off according to the menu cost model, and no such relationship in the new classical model, provides a basis for choosing between them. The authors' basic procedure is to estimate short-run output-inflation trade-offs for 43 industrialized countries and then to explain the differences in the estimated coefficients across countries by differences in the average values of the variables that the theories predict should be important. The trade-off for each country is estimated by relating real GNP to its lagged value, time, and the change in nominal GNP, using annual data for the period 1948–86 and for the subperiods obtained by splitting the sample in 1972. The estimated coefficients vary widely. On average, approximately one-quarter of a shock to nominal GNP shows up in output in the same year; the standard deviation of the fraction is 0.27. The fraction for the United States is 0.67, showing a far larger than average effect on real output from shocks to nominal demand. While the coefficient for the United States is essentially the same for the two subsamples, the correlation between the subperiod estimates for all countries is only 0.36, indicating there can be considerable variation in the trade-off over time.

Simple plots of the trade-off coefficients against mean inflation and the standard deviation of nominal GNP growth show strong evidence of a negative and nonlinear relation in both cases. However, since mean inflation and the standard deviation of nominal growth are highly correlated, the authors turn to multiple regressions to determine whether

these apparent relationships should be attributed to one or the other or both of the variables. These regressions show that the effect should be credited to mean inflation rather than demand variability; the partial effect of demand variability is actually estimated to have the wrong sign. The estimates indicate that increases in the mean rate of inflation have substantial effects on the output-inflation trade-off. For example, an increase in the mean inflation rate from 5 percent to 10 percent reduces the fraction of a nominal shock going into output in the first year by 0.22.

The authors examine the robustness of this basic result in a variety of ways. They report regressions relating the change in the trade-off coefficient between the subsamples to the change in the mean level of inflation and the change in demand variability, which in effect eliminates any fixed country effects. They also investigate the possibility that supply shocks are giving rise to spurious results, or that the particular sample of countries is important. None of these appears to make a significant qualitative difference to their results.

The authors conclude that the evidence supports new Keynesian, not new classical, theories of the output-inflation trade-off. The correlation between the variance of demand shocks and the trade-off that previous research had cited in support of new classical models appears to reflect the correlation of the variance of shocks with the average level of inflation, rather than a structural relationship. Although the models have quite different predictions about some policy actions—for example, fully anticipated monetary policy can have an effect on real output in the authors' model—some of their model's predictions would have a familiar ring to a new classicist. In particular, the response of output to nominal shocks is not invariant to the rate of monetary growth, a result predicted by the Lucas critique. A permanent increase in the rate of monetary expansion not only eventually shows up entirely in higher inflation, but reduces the response of output to transitory monetary actions.

DEBT LEVELS of all kinds have been rising in the U.S. economy in recent years, causing concern among politicians, economists, and financial analysts. To the casual observer, the debt arising from federal deficits, the debt to foreigners arising from trade deficits, the developing country debt held by banks, the nonperforming debt held by thrift institutions, and the growing indebtedness of consumers and nonfinancial businesses may appear simply as one general debt problem. But these different

types of debt should not be lumped together. They differ in the magnitude and incidence of the debt burden they pose, in the risk of default and the likelihood of triggering a broader financial crisis, and in the role that government policy can and should play. Hence they need to be analyzed separately. The economic consequences of deficits in the federal budget and foreign trade, the developing country debts of banks, and the financial problems of thrifts have all been investigated in previous issues of the *Brookings Papers*. In the second paper of the present issue, Ben S. Bernanke and John Y. Campbell examine the debt levels of nonfinancial firms in the United States and assess the potential risks arising from firms' increased debt leverage.

Bernanke and Campbell show that the rise in corporate debt relative to income has been a recent phenomenon. During the 1970s, the growth of total corporate debt just kept pace with the growth of nominal GNP. Not until 1984 did this relation change noticeably, with corporate debt rising 18 percentage points more than nominal GNP during 1984–86. The authors show, for a sample of firms in the COMPUSTAT data bank, that the growth of debt during 1984–86 corresponded with a change from past patterns of net equity issues by corporations. Until 1984, the firms in this sample issued small net amounts of equity each year; but during 1984–86, they repurchased large amounts of equity, mostly in connection with takeovers and financial reorganizations. Hence recent years have seen a substantial substitution of debt for equity.

What are the social consequences of this greater leverage? Greater leverage may make it more difficult for a firm to raise funds in the future at a time of financial stringency and, indeed, increases the risk that the firm will go bankrupt. But these costs are obviously borne directly by individual firms and, if management of the firms is rational, should be part of the calculation in taking on the debt in the first place. However, the authors point out that there are also negative externalities to the macroeconomy as a result of increased firm leverage. If firms with high leverage get into trouble because of a recession, the added difficulty and cost of raising funds would adversely affect their real investment, tending to worsen the recession. Bankruptcy may even have contagion effects if, as some firms go bankrupt, others find lenders less willing to provide additional finance than they otherwise would be. In addition to the externalities felt within the corporate sector itself, widespread financial distress could reduce confidence throughout the economy, contributing

to a recession by reducing the demand of consumers and firms in other sectors.

In an evaluation of the importance of issues such as bankruptcy risk, aggregate statistics, which reflect the condition of the median firm, are less revealing than statistics that reflect the condition of firms in the tail of the distribution whose exposure to risk is greatest. To get at the tail of the distribution of firms, the authors rank the firms in the COMPUSTAT data bank from the most risky on down. They focus on two measures of the vulnerability of firms to financial stress: the ratio of the market value of debt to the market value of assets and the ratio of interest obligations to cash flow. The debt-to-asset ratio indicates whether the present value of expected future earnings is sufficient to cover the present value of contractual debt payments. In a world of well-functioning capital markets, a firm with a low debt-asset ratio should be able to finance any short-term cash shortfall rather than go bankrupt. The ratio of interest to cash flow indicates whether firms are able to meet their interest obligations without additional financing. With imperfect capital markets, a firm with a high ratio may find new financing expensive and could even go bankrupt although future earnings would cover debt obligations. This second measure also lends itself to projecting the effects of deteriorating economic conditions.

For their analysis, the authors adjust the book value of the debt of the firms in their sample to market values, and relate these to the market value of the firm. Because the stock market rose during the same years that companies increased their debt, the authors find no trend in debt-asset ratios during recent years, either for the average of their firms or for firms in any part of the distribution. Although they regard debt-asset ratios as the best indicator of the solvency of firms, they note two reasons why such ratios may not provide an adequate picture of the financial vulnerability of firms. First, the stock market exhibits much unexplainable variation and may not be a reliable measure of fundamental value. Second, even correctly priced stocks will more nearly reflect average rather than worst-case scenarios.

Bernanke and Campbell find more evidence of change in their second indicator of financial soundness, which reflects liquidity: the ratio of interest paid by the firm to its cash flow before depreciation, taxes, and interest. They find that this interest-coverage measure has deteriorated during the 1980s, particularly for the 90th and 95th percentile of firms in

their sample. For example, for 95th percentile firms this ratio varied between 0.5 and 0.9 during 1971–79, but has exceeded 1.28 every year since 1980. In every year since 1981, these firms have looked as vulnerable as Chrysler did in 1978. When the authors use a three-year moving average of cash flow to smooth out temporary disturbances to a firm's performance, the interest-coverage ratios are less erratic and the pattern through time somewhat clearer. For firms in both the 90th and 95th percentile, these moving average interest-coverage ratios rise noticeably after 1983. For 95th percentile firms, the ratio never exceeded 0.85 before 1981, was 1.1 in 1981, and grew rapidly to 4.7 by the end of the sample.

The main danger from increased corporate leverage is an event, such as a deep recession or a money crunch, that significantly lowers corporate earnings relative to interest obligations. To evaluate this danger, Bernanke and Campbell simulate the condition of firms in a two-year recession beginning in 1987 that is similar to the recession of 1973–74 or 1981–82. They do this by computing, for each firm in their sample that existed in 1972 or 1980, the percentage changes in market value of debt and equity, cash flow, and interest expense in the two succeeding years. They find that net asset ratios would deteriorate most under conditions simulating the 1973–74 recession, when the stock market fell very sharply. Under these conditions, debt-asset ratios would approach 1.0 for firms as low as the 75th percentile and exceed 1.0 for firms above the 90th percentile. Simulations of the 1981–82 recession show a noticeably smaller effect on debt-asset ratios, but interest coverage deteriorates much more. By the second year, firms at the 90th percentile would reach interest-coverage ratios as bad as those of the Chrysler Corporation during 1979–80. The simulations also show how different the exposed firms are compared with the median firms, whose interest coverage would rise only from 0.18 to 0.25 in the recession. Even firms at the 75th percentile would still have interest-to-cash-flow ratios less than 0.5.

Based on these recession simulations, the authors warn that, by postwar standards, an unprecedented proportion of firms could end up financially distressed in any major new recession. However, they also note that there is no precise estimate of what level of their measures of financial distress would actually result in bankruptcy or a serious disruption in a firm's ability to raise capital. Comparison with historical experiences, such as those of the Chrysler Corporation, provides some

evidence. But the acceptable levels for these ratios may conceivably have changed, making any particular value, either of debt to assets or interest expenses to cash flow, less risky today than in the past.

One argument is that junk bonds are held primarily by insiders who hold equity stakes along with junk bonds and who thus could be counted on to renegotiate rather than force bankruptcy. The authors do not find this argument convincing for several reasons. They note that although junk bonds constitute about 20 percent of corporate bonds outstanding in 1986, they account for well under 10 percent of total nonfinancial corporate debt, including bank loans and short-term debt. Junk bond issues have increased mainly as a replacement for bank loans, and the authors see them as part of the general trend toward "securitization" of financial instruments. They also observe that junk bonds are widely held by mutual funds and other institutional investors, not just insiders, making it difficult to avoid bankruptcy through negotiation and the provision of new funds.

Another argument is that overfunding of corporate pension funds has given corporations flexibility to take on more debt than in the past. But the authors note that pension fund surpluses have been eliminated in recent years by the decline in interest rates that increased the present value of pension liabilities. Thus they regard it as unlikely that changes in pension fund balances could rationalize higher interest burdens.

Not being persuaded by arguments that high levels of debt represent substantially less exposure to financial crises than they have historically, the authors conclude with an expression of concern. If the strong earnings growth implicit in current stock prices fails to materialize, in particular if a major recession should occur, the financial health of the corporate sector could be seriously threatened.

BECAUSE OF THE IMPORTANCE of corporate investment, both to aggregate demand and to the productive capacity of the U.S. economy, the study of its determinants has long been a major focus of research and debate. Although cash flow and investment are highly correlated, classical, neoclassical, and Keynesian theories all focus on the relationship between the marginal product of capital and "the" interest rate, not on cash flow, to explain investment. Perhaps this common theoretical ground, and the absence of a theoretically attractive alternative, helps explain economists' continued emphasis on interest rates, despite a lack

of any clear connection between interest rates and investment in the data. In the third paper of this issue, Steven M. Fazzari, R. Glenn Hubbard, and Bruce C. Petersen reexamine the theoretical rationale and empirical support for using the internal cash flow of firms in explaining investment.

Under the standard neoclassical assumption that capital markets are perfect, there should be no link between a firm's financial structure and its real investment spending; internal and external finance should be perfect substitutes. However, Fazzari, Hubbard, and Petersen point out that reasons abound why internal finance should be less expensive than external finance. Among the most prominent are transactions costs, tax advantages, agency costs, real bankruptcy costs, and asymmetric information between borrowers and lenders about the prospects of the firm. The transactions costs associated with new share issues, for example underwriting and registration fees, can be substantial for small offerings. Similarly, the tax system, which taxes gains only on realization and has historically taxed realizations at a lower rate than dividends, makes the cost of retained earnings lower than the cost of external equity finance. While the authors believe that such effects may be important, they emphasize the cost advantage of internal funds that arises because of asymmetric information between managers and potential new investors or creditors.

That managers know more than the market does about the profitability of their own investment projects introduces a "lemons" problem of the kind first analyzed by George Akerlof. Managers with highly profitable investment opportunities will not want to sell new equity to less well-informed buyers, because buyers require a lemons premium. This basic idea provides a theoretically rigorous argument for a financing hierarchy, with internal funds preferred to new equity issues. Asymmetric information can introduce a similar distortion in the issue of debt, with low-risk borrowers dropping out of the market, and provides a reason why banks and other lenders find it profitable to ration credit. Asymmetric information also brings about agency costs, which drive a wedge between the interests of managers and bondholders and provide incentives for "second-best" restrictions on debt instruments. Such restrictions, for example stipulating a target debt-equity ratio or imposing working capital requirements, may protect the lender from some costly managerial actions, but the restrictions themselves are not costless. They thus help

explain why the cost of borrowing may exceed the opportunity cost of using internal funds for investment. The authors observe that such restrictions are likely to be most important at precisely those times when the need for external finance is most acute.

Fazzari, Hubbard, and Petersen modify the standard model of financial and investment decisions, which is based on maximizing the present value of the firm, to incorporate a difference between the cost of internal and external finance reflecting asymmetric information. Not surprisingly, the model generates a financing hierarchy and also has a number of implications for q—the ratio of a firm's market value to the replacement cost of its capital—and its relationship to investment. First, the model demonstrates that, to induce a given level of investment, q will have to be higher for firms about which the market has limited information than for firms about which the market is well informed. Second, for low levels of q, investment should be related to internal funds and only loosely related to q, whereas for high levels of q the association of q and investment should be much closer.

In contrast with most recent studies of investment demand, Fazzari, Hubbard, and Petersen explicitly allow for a heterogeneous population of firms. They expect liquidity constraints to affect mature firms, on average, less than younger firms. Mature firms' investment needs are less likely to exceed their internal cash flow, and they are better known in financial markets so that their cost of credit is less likely to include a lemons premium. Hence, they believe it is inappropriate to assume a representative firm and to test models of investment by comparing how well they explain a representative firm's behavior. In their view, there are good theoretical reasons for believing that some firms may be liquidity constrained while others are not, with aggregate behavior reflecting a blend of the behavior of firms of each type.

The authors examine the empirical importance of financial constraints using a panel of manufacturing firms drawn from the Value Line data base. The authors divide the firms into three different classes that, a priori, they expect to have different costs of external finance. Firms in class 1 have ratios of dividends to income of less than 0.1 for at least 10 of the 15 years in the sample; firms in class 2 have ratios between 0.1 and 0.2 for at least 10 of the 15 years; and all other firms are in class 3. The differences in dividend behavior are substantial. Class 1 firms retained on average 94 percent of their income and paid dividends in only one-

third of the years of the sample period, while class 3 firms retained on average less than 60 percent of their earnings and paid dividends in 98 percent of the years. Class 1 firms have less than one-tenth the capital stock of class 3 firms at the beginning of the sample period, and much higher growth of the capital stock over the period. Class 1 firms had average sales growth of 13.7 percent, compared with 8.7 percent and 4.6 percent for class 2 and 3 firms, respectively. Differences in the relationship between investment and cash flow across the three types of firms are also substantial. The average correlation of cash flow and investment, each detrended, is 0.9 for class 1 firms and only 0.2 for class 3 firms. At the same time that the investment of class 1 firms appears more tightly constrained by cash flow, class 1 firms are much more likely than class 2 or 3 firms to issue new equity and are more heavily in debt. Debt does not appear to be used to smooth investment relative to earnings. Rather, for all classes of firms, and most significantly for class 1 firms, changes in debt are positively correlated with earnings.

Since the authors' hypothesis relates to the elasticity of the supply of funds for investment in physical capital, not to the demand for funds, it can be tested together with any one of a variety of theories about the determinants of investment demand. The authors examine the importance of financial constraints by testing the significance of cash flow in three different models of investment: a q model using tax-adjusted market value relative to replacement cost ("Q") as an explanatory variable, a neoclassical model using a measure of output and the cost of capital, and a sales accelerator model. The results for the three models are quite similar and the authors focus their attention on the q model. Although the coefficient for cash flow is significant for all classes of firms, the coefficient for class 1 firms is significantly greater than that for the other two classes. Depending on the sample period, approximately one-half to two-thirds of an extra dollar of cash flow of the average class 1 firm goes into investment. In contrast, the fraction is approximately one-third for class 2 firms and one-quarter or less for class 3 firms. The cash flow coefficient for class 1 firms gradually declines as the sample period is extended from 1975 to 1984, consistent with the authors' hypothesis that as class 1 firms mature over the sample period, cash flow should become less important. The behavior of the coefficient on Q itself is also consistent with the authors' hypothesis. It does not perform well for class 1 firms but is positive and significant for firms in classes 2 and 3.

A standard caveat in interpreting cash flow regressions is that current cash flow may simply be a proxy for future profitability that is not well captured by current profits, sales, or output. However, in the authors' regressions, Q presumably incorporates the market's expectations of future profitability. The success of cash flow in the Q equations, and later in the sales accelerator models even when Q is included, therefore suggests that cash flow in fact reflects financial constraints. The authors try many different specifications of the q model, using lagged values of Q as an instrument for Q, taking out fixed year and firm effects, running the equations in first or second differences, and using instruments for cash flow itself in an effort to deal with its possible endogeneity. The results from the basic q specification are robust to all these changes.

Fazzari, Hubbard, and Petersen conclude that financial constraints on investment are important for many firms. But how important are such constraints to aggregate investment and in particular to its cyclical fluctuation? The authors calculate a lower bound on such effects, attributing to financial constraints only the portion of the cash flow coefficients of class 1 and 2 firms that is in excess of the coefficient for class 3 firms, and assuming that the proportion of financially constrained firms in the entire economy is no greater than the proportion of such firms in the Value Line survey. With these conservative assumptions, they show that a one standard deviation change in the cash flow of financially constrained firms explains about 13 percent of the standard deviation in aggregate investment.

The authors' findings have important implications for the effect of tax policies on investment. For financially constrained firms, the average, rather than marginal, effects of changes in depreciation allowances, tax rates, or investment tax credits are important. The cost of capital effects of tax policies are usually estimated to be small; cash flow effects may therefore be the most important channel by which policy affects aggregate investment.

KEYNESIAN MODELS, the dominant macroeconomic paradigm in the 20 years after World War II, have been under challenge since the late 1960s. The acceleration of inflation during the Vietnam War and the onset of stagflation during the 1970s—economic setbacks that, to the public, seemed to be at odds with the optimistic predictions of Keynesian economics—caused the public to lose confidence in the standard policy prescriptions identified as Keynesian. Within the profession, the chal-

lenge to Keynesian economics arose not so much from these events but from the lack of rigorous microeconomic foundations in Keynesian models and from the emergence of the new classical models, which built on the dominant microeconomic paradigm of rational agents and competitive market clearing. In recent years, partially stimulated by the failure of the new classical model to explain the experience of the late 1970s and 1980s, there has been a resurgence of interest in Keynesian models, but with a new emphasis on microeconomic underpinnings for the crucial Keynesian features. In the first and third papers in this issue, two features of such new Keynesian models are explored in detail. Ball, Mankiw, and Romer review the menu cost theory of price and wage rigidity and test its macroeconomic implications. Fazzari, Hubbard, and Petersen examine the theoretical rationale for financial constraints on firms' investment and test that theory using firm data. In the fourth paper of this issue, Bruce C. Greenwald and Joseph E. Stiglitz, who unabashedly favor the new Keynesian models, examine the ability of the three major competitors—new classical, traditional Keynesian, and new Keynesian theories—to explain important stylized facts about the performance of industrial economies.

Greenwald and Stiglitz search for stylized facts with which to confront these theories by examining time series characteristics of macroeconomic variables that they regard as key. They examine data from several countries, rather than just from the United States, arguing that basic theories should be tested against the performance of industrial economies in general. They consider data relating to three markets of central importance: goods markets, capital markets, and labor markets.

In goods markets, the authors show that deviations of output from trend (in all cases, the trend is allowed to change frequently) are serially correlated, although there is no clear correlation in deviations of changes in output from trend. Fluctuations in output or in changes in output are ubiquitous and similar in size across countries, and also similar in the first and second halves of U.S. postwar experience. Price inflation is more persistent than are output changes, with deviations of inflation rates from trend positively serially correlated in all countries.

Lacking reliable data on real interest rates, the authors characterize capital markets by trends in real investment. For all countries and time periods in their sample, they show that total investment fluctuations are about three times as large as output fluctuations, measured either as

deviations of levels or first differences from trend. Fluctuations in residential construction investment are exceptionally large. Investment fluctuates more relative to output in some countries than others, with investment in Japan being most stable.

In the labor market, the authors show that correlations between deviations from trend in real wages and output are very low, with some small positive correlations in some countries (procyclical fluctuations) when the consumer price index is used to deflate wages, and a mix of small positive and negative correlations (countercyclical fluctuations) when producer prices are used. Both employment and hours are procyclical, but their correlations with output vary considerably across countries. The Depression led to some departures from these general relations. Real wages measured by consumer prices rose significantly when food prices fell, producing a negative correlation with changes in output and employment. The authors also show that productivity growth is procyclical everywhere. Finally, they show that except in Japan, where historical unemployment data are unreliable, changes in unemployment exhibit persistent departures from trend.

Greenwald and Stiglitz describe simple versions of the three models of business-cycle behavior in order to compare their predictions against these stylized facts. Their version of the real business-cycle model has two basic elements that they believe are representative of most such models: a production function in which the capital stock, labor, and technological shocks determine output, and a representative consumer who, in each period, maximizes expected utility over an infinite horizon by choosing labor supply, consumption, and, implicitly, investment, using all available information about the capital stock and technology. Behavior is competitive.

The real business-cycle model, in spirit and structure, is an extension of the usual competitive general equilibrium model and shares its basic properties. Greenwald and Stiglitz believe that the most important of these is the general tendency for competitive markets to dampen the effect of external disturbances. In contrast to Keynesian models, in which shocks can be amplified and, indeed, fluctuations in output can be largely endogenous, in real business-cycle models the source of economywide output fluctuations at business-cycle frequencies is technological, with the magnitude and persistence of cycles depending on the nature of those shocks.

Without any restrictions on the distribution of such shocks, it is clearly possible to match quite precisely the empirical properties of real output fluctuations described above. But Greenwald and Stiglitz argue that it is difficult to reconcile the observed behavior of output with the view that it is driven by technological shocks. They argue, for one thing, that while technology advances may be quite uncertain, technology should not regress. It is therefore difficult for real business-cycle models to explain those instances in which output and productivity fell during the postwar period, to say nothing of the Great Depression. Such models also have difficulty accounting for the apparent size and cross-country correlations of shocks. The authors calculate that if aggregate output were composed of the output of 50 independent industries, the required standard deviation of industry shocks per quarter would be 5 percent, which seems implausibly high relative to estimates of fluctuations in industry productivity. They further observe that if technology is commonly available, the output of industrial economies should be closely correlated, and they note that many such correlations are negative and few strongly positive. Similarly, output correlations ought to be higher for a given industry across countries than across industries within a given country. Yet the available data suggest the opposite—national factors seem more important than industry factors.

Because technological shocks increase the demand for capital, real business cycles can generate fluctuations in investment that are much larger than fluctuations in output, in accord with the authors' stylized facts. However, the authors note an awkwardness for the model in that recessions should, but do not, lead to negative net investment. The model has mixed success in explaining the pattern of fluctuations across categories of investment. The authors show that the model can explain the relative variability of investment in business structures and durable equipment, but substantially underpredicts the variability of residential construction and inventory investment.

The labor market poses the most difficult challenge for the authors' real business-cycle model. The model assumes that individuals are always on their labor supply schedules so that fluctuations in employment reflect the response of workers to real wage rates. The authors first note that the highly elastic relationship between wages and hours evident in the aggregate data seems inconsistent with cross-sectional studies that suggest an inelastic labor supply. Again, the Depression provides a

dramatic test: it is hard to believe that the 22 percent decline in nonfarm employment between 1929 and 1932 left workers on their supply curve. Indeed, under what the authors believe to be plausible assumptions about workers' preferences, the rise in real consumption wages during the Depression should have led to an increase in hours worked. Even less plausible to the authors is the notion that utility-maximizing workers would choose to cease working entirely. Again, observed behavior does not appear to be consistent with the view that workers are on their supply curves.

Greenwald and Stiglitz also find an inadequate correspondence between facts and theory when they examine traditional Keynesian models. They characterize cycles in such models as arising from sticky nominal wages in the face of shocks to nominal demand. Employment is determined by firms' demand for labor so that workers are not on any well-defined labor supply schedule. The authors identify this wage stickiness with overlapping wage contracts that can perpetuate the labor market disequilibrium, though, they believe, for only a limited time. Although many exponents of Keynesian theory would not subscribe to such a narrow explanation for wage stickiness, from it the authors argue that traditional Keynesian theories can explain only brief departures of output from trend. Thus they regard the extreme persistence in output fluctuations apparent in data for the U.S. economy as evidence against this model. They also find other stylized facts inconsistent with their version of the traditional Keynesian model. Demand shocks are the dominant source of cycles in this model. Because the authors believe that differences in degree of openness and institutional arrangements should make some economies much more stable than others, they regard the broad similarity in output fluctuations in different economies as evidence against the traditional Keynesian model. Furthermore, their model offers no clear explanation why prices as well as wages should be rigid, and so offers no explanation for the relative constancy of real wages over the cycle.

The authors identify the investment accelerator with the traditional Keynesian model. But because they regard that model as predicting only temporary departures of output from trend, they argue that it should predict only small accelerator effects that would not explain the observed data on investment fluctuations. Similarly, their traditional Keynesian model does not incorporate a financial structure that would provide

explanations rooted in financing for fluctuations in residential construction. In their model, the large fluctuations observed in residential construction are hard to explain because the demand for housing depends on lifetime, rather than current, income, and therefore should not fluctuate wildly.

The new Keynesian theories rationalize and elaborate the traditional Keynesian assumptions in a number of ways: new theories of search unemployment, implicit contracts, and efficiency wages provide a more detailed description of the labor market and a more rigorous explanation of wage rigidities; new theories of menu costs and imperfect competition explain price rigidities in terms of rational behavior; credit and equity rationing are rationalized by asymmetric information that results in adverse selection and moral hazard.

Greenwald and Stiglitz use a model built around equity-rationing constraints—limited recourse to sales of new equity as a means of raising external capital—and efficiency wages that they believe illustrates the distinctive implications of the new Keynesian models. In their model, firms pay factors of production at fixed rates. Output arrives in the future and is of uncertain value. Equity is changed primarily by retained earnings, because the firm's access to public equity markets is limited. Hence a decision to increase production carries a risk of worsening the firm's equity position and indeed increases the chances of bankruptcy. The smaller the firm's equity base relative to its debt and other fixed obligations, the more serious that risk. Equity becomes a state variable influencing the firm's employment, output, and investment decisions.

An adverse demand or supply shock leads to an immediate deterioration of the firm's balance sheet, because the value of assets related to production will immediately fall, whereas the commitments incurred in acquiring these assets will not fall commensurately. The firm's balance sheet will only gradually be restored through retained earnings, and hence reductions in output may be extremely persistent. Furthermore, output of firms should be positively correlated because the output decisions of one firm affect the cash flow of others. Both these predictions of the model accord with the authors' stylized facts. The interaction between firms also provides a mechanism by which shocks can be amplified by the system, explaining aggregate fluctuations of the observed magnitude with disturbances of a more plausible size than required by the real business-cycle models. In the Greenwald and Stiglitz

new Keynesian model, investment fluctuations are large relative to output because deferring investment is one of the least costly ways of reducing risk as a firm's financial position deteriorates.

The assumption that firms pay efficiency wages explains several stylized facts about wages and employment. According to efficiency wage theory, higher real wages improve labor productivity by eliciting greater effort from workers, attracting a high-quality labor supply, and reducing turnover. For this reason, firms may choose real wage rates that are too high to clear the labor market, so that involuntary unemployment may be a feature of equilibrium. Because workers may be off their supply curve, the model can explain the apparently elastic relationship between real wages and employment observed in the real world. It reflects the way firms optimally adjust efficiency wages and employment to changes in demand. Because firms believe that increases in demand are associated with a reduction of risk to the marginal product of labor, the risk-adjusted marginal product may actually rise, rather than decline, during an expansion, warranting a rise in the real wage, in accordance with the facts. However, the authors note that the model cannot explain the observed procyclical movements in measured productivity.

Greenwald and Stiglitz's overall assessment is that neither the real business-cycle model nor the traditional Keynesian model is particularly successful in explaining the stylized facts in any of the markets investigated. They find their new Keynesian approach more successful. However, they conclude that none of the models, at least in the simple variants they present, successfully explains all of the salient features of the business cycle.

UNEMPLOYMENT in Europe rose from 3 percent in the 1970s to near 10 percent during the first half of the 1980s and has stayed on this high plateau ever since. In the first report of this issue, Robert J. Gordon considers various explanations for high European unemployment and provides an empirical analysis supporting his view that unemployment would respond to an expansion of demand with limited, rather than accelerating, inflation. Although nothing comparable to the Great Depression of the 1930s has devastated Europe in this decade, Gordon sees some similarities between the macroeconomic conditions of Europe in the mid-1980s and the United States in 1939. In both cases, a range of supply-side conditions existed that many observers today believe would

inhibit expansion—high real wages, union militancy, and a declining capital stock. Noting that the United States expanded smartly once demands stemming from war in Europe stimulated the economy, he suggests that a demand expansion would be similarly successful in Europe today, though it would entail some increase in inflation rates.

Before turning to his empirical analysis, Gordon examines the evidence for several views of European unemployment that are more pessimistic. The first is a group of structural explanations that blame government regulation and the welfare state for Europe's high unemployment. In assessing the structuralist explanation, Gordon cites studies showing that the effect of unemployment benefits is far too small to account for much unemployment, and that European replacement rates—the proportion of working income that is replaced by unemployment benefits—have generally fallen rather than increased since 1980, while unemployment has risen substantially. In Germany, for example, the replacement rate appears to have fallen from 89 percent in 1970 to 26 percent in 1984. Gordon concedes that taxes as a share of GDP have risen continuously in Europe during the period of rising unemployment, and that these increases could, in principle, have reduced work incentives. But he observes that the rise in taxes in Japan is almost as large and has not interfered with employment expansion there. Gordon also questions whether employment security arrangements in Europe are important, citing work that shows they are, on balance, not noticeably more restrictive than those in the United States. He reports that minimum wages have declined in most European countries and do not exist in Germany and Italy, two countries with relatively high unemployment. He concedes that some regulations may have inhibited expansion; for example, restrictions on entry may well have slowed the growth of convenience stores in the service sector in Europe. However, he is not persuaded that welfare state policies and regulation are responsible for increasing structural unemployment. Gordon also rejects the idea that high real wages inhibit greater expansion of employment in Europe. He shows that real wages and the wage gap for European countries have generally declined since the late 1970s rather than rising as the real wage explanations for European unemployment would require.

Because inflation in Europe has stabilized at a plateau of high unemployment, many analysts have concluded that the European NAIRU—the unemployment rate at which inflation neither accelerates

nor decelerates—has moved up and corresponds, roughly, to the present actual unemployment rate. The implication is that demand expansion cannot reduce unemployment without causing accelerating inflation. Many analysts associate such a hypothesized rise in the NAIRU with structural causes of unemployment of a microeconomic nature. Because Gordon does not find persuasive evidence for any large permanent effects in the usual list of microeconomic suspects, he prefers a different explanation for the observed stagflation, one based on the hysteresis hypothesis. In the short run, that hypothesis offers an apparent NAIRU that in fact is endogenous, tracking the actual level of unemployment with a lag. The distinctive feature of hysteresis explanations is that a demand expansion that reduced actual unemployment would, with a lag, reduce the apparent NAIRU with it.

Hysteresis explanations do not necessarily imply that inflation would not quicken in the process of such an expansion. Indeed, Gordon presents a formal model illustrating the point. Suppose inflation is related to its own past value with an elasticity near 1.0 and to the gap between actual unemployment and lagged unemployment. In such a model, reducing current unemployment through demand expansion adds to inflation in the same period, but inflation does not increase further if unemployment is held constant because the gap between current and lagged unemployment disappears. The original increase in inflation is perpetuated through the lagged inflation term. The NAIRU is simply lagged unemployment and has no structural connection with accelerating inflation. If the elasticity between lagged inflation and current inflation is exactly 1.0, then, from any initial position in which the NAIRU and unemployment are equal and inflation is steady, unemployment can be reduced to a new permanent level and inflation will, in the process, rise to a new permanent level. The force of the NAIRU idea therefore depends critically on how rapidly, if at all, it responds to the level of unemployment itself.

The three principal explanations that have been offered for why the apparent NAIRU has risen along with the actual unemployment rate are the disappearance of physical capital, the decay of human capital, and the distinction between insider and outsider members of the work force. While some of these are hard to measure, Gordon observes that investment has not been low in Europe and that Europe appears to be amply endowed with capital. Furthermore, if any of the hysteresis explanations is taken to imply that policy has to accept the current high levels of

unemployment, it must be that the lags involved in bringing down the NAIRU would be sufficiently long that the inflationary cost of doing so would be unacceptable. As evidence against this, Gordon recalls past occasions when economies have achieved a rapid growth in output and employment after long periods of high unemployment when capital and skill deterioration must have been even more severe than in Europe today, and when insider-outsider relations should have been equally important.

Gordon prefers an explanation of recent European experience that places less emphasis on these relatively slow adjustments of the environment for wage setting to actual unemployment and more emphasis on rate-of-change effects. He observes that inflation equations embodying the hysteresis explanation are formally indistinguishable from explanations in which short-run inflation responds to changes in unemployment rather than to its level. He hypothesizes that the effect on inflation of changing unemployment or, equivalently, of a changing output gap, is a structural feature of economies; it arises because, when unemployment is rising, the employed fear for their jobs and moderate their wage demands, while, when unemployment falls, employers have to raise wages to attract new workers and keep old ones from quitting. By contrast, he reasons that a stable level of unemployment exerts little pressure on wages, either up or down, as long as the level of wages is consistent with labor productivity.

Gordon estimates wage and price equations to test his explanation of European experience. Changes in either the GDP deflator or unit labor costs are explained by current and two lagged values of the output gap, the current change in relative import prices, a dummy variable for special wage-push effects in 1968–70, and two lags of the dependent variable. In support of his general view, he finds that present inflation has an elasticity of about 1.0 with past inflation and that changes in the output gap lagged a year, rather than the level of the gap, are significant. He also reports that results for the United States during the interwar period—1922–39—show that changes in output gaps, rather than their levels, are important for the rate of inflation. However, the interwar U.S. equations show an elasticity between present and past inflation of only 0.4 rather than the 1.0 found for present-day Europe.

Gordon shows the implication of his inflation equations by simulating a future demand expansion in Europe. When the output gap, which

Gordon estimates to have been 3.4 percent in 1986, is maintained over the full 1987–96 decade, thus keeping output growing at its trend rate, inflation declines gradually from 4.2 percent to 3.1 percent by 1996. When the output gap is reduced steadily over the five years 1987–91 and held at zero thereafter, inflation rises to approximately 6 percent by 1993. When output is allowed to expand twice as fast, resulting in a negative output gap of 3.4 percent after 1991, inflation rises to 8.9 percent by 1996. In all cases, the rate of inflation is essentially constant after the initial period of adjustment. Hence, according to Gordon's results, although there is an inflation cost to expanding output and reducing unemployment, the output gains can be permanent without leading to an accelerating inflation. The trade-off between output and inflation is thus much better than many observers have believed.

BEGINNING IN 1982, average annual wage increases in the U.S. economy slowed sharply. They have thus far remained moderate throughout the economic expansion that reduced the aggregate unemployment rate from over 10 percent at the end of 1982 to 5.5 percent in the second quarter of 1988. Over this period, increases in hourly earnings in the private nonfarm economy averaged less than 3 percent a year, lower than in any period since the early 1960s. In the final report of this issue, Wayne Vroman and John M. Abowd examine this experience with particular attention to the effect of foreign competition on wage moderation.

The authors first use conventional aggregate inflation equations that explain the hourly earnings index by lagged consumer price inflation, unemployment, and lagged import price inflation. Most of the explanation comes from lagged consumer price inflation, with a coefficient of approximately 0.75. They find that the unemployment rate for men aged 25–54, which has been high relative to the total unemployment rate in the 1980s, explains wage movements better than the total unemployment rate both historically and during the slowdown of wage inflation in the 1980s. They also find a significant, though not large, positive effect from import price inflation. Although the preferred forms of these equations predict wage changes of the 1980s reasonably well, even when those years lie outside the estimation period, the equations overpredict wage increases by almost a percentage point in 1986 and 1987.

Detailed data on major collective bargaining settlements in manufacturing industries provide the authors further evidence on wage behavior.

Their data set, with information on 2,700 individual agreements from 1959 though 1984, allows them to include many variables not available at the aggregate level that may influence wage settlements at the bargaining table. The authors allow for three types of explanatory variables: aggregate labor market conditions, proxied by the unemployment rate for men aged 25 to 54 and by the most recent settlements in the auto and steel industries; both expectations and realizations of price inflation, with distinctions between contracts with and without cost-of-living adjustments (COLAs); and industry-level indicators of product market conditions, including import penetration ratios and the change in industry exports, as well as changes in the price and volume of domestic shipments.

Projecting the years 1980 to 1984 from equations estimated through 1979, Vroman and Abowd find a predominance of overpredictions, although the errors are no larger than typical annual errors during the sample period. Over the projection years, wage settlements declined by about 6 percentage points. According to the estimated equations, about 2 percentage points of decline are explained by the decline in expected inflation, measured by lagged values of the consumer prices. Another 1 percent is explained by catch-up to inflation surprises during the previous contract; 2 percent by slowdowns in the auto and steel settlements; and 0.5 percent by the level and change in the unemployment rate. Because the auto and steel settlements, as well as expected and realized inflation, all reflect the effects of unemployment, the total effect of unemployment in the slowdown is greater than this decomposition of effects indicates.

The authors' estimates suggest that there are special effects from import competition but not from rising exports. An added dollar of import penetration is estimated to have an effect several times greater than a dollar of lost domestic consumption. However, the effects are not quantitatively large. In a simulation with imports growing one standard deviation faster than domestic consumption, after six years annual wage settlements are considerably less than 1 percentage point smaller in all but three of twenty industries analyzed.

Vroman and Abowd observe that, since 1982, union wage increases have slowed more than nonunion increases. They also note that the relative importance of major unions in total private employment has continued a decline that had started in the 1960s. This union employment decline is not due mainly to changes in the relative importance of heavily

unionized industries; between 1978 and 1987 it appears in virtually all two-digit Standard Industrial Classification code industries. In a time series regression aggregating across all highly unionized industries, the authors find that the decline in the major union percentage of employment is explainable by a trend that steepens after 1976, the unemployment rate, and real imports as a share of GNP. Using this equation, they find that the rise in the import share accounted for an estimated 41 percent of the total decline in the major union percentage between 1968 and 1987. Regressions across industries gave a significant but smaller estimated effect, with import penetration accounting for 10–20 percent of the decline. From such necessarily imprecise estimates, Vroman and Abowd conclude that international competition has been important, but can explain only a part of the decline in major union employment.

Whatever the source of the decline in unions' economic power, Vroman and Abowd observe that it has tilted unions' objectives toward greater job security relative to pay gains. The 1980s saw the spread of concession bargains, lump sum payments in place of permanent wage increases, two-tier pay arrangements, and reduced COLA protection as well as lower negotiated wage gains. Vroman and Abowd believe that such sweeping changes in union and management practices will not be easily reversed, and so expect continued moderation in union wage increases compared with previous periods of high employment.