Editors' Summary

THE BROOKINGS PANEL on Economic Activity held its sixty-ninth conference in Washington, D.C., on March 30 and 31, 2000. This issue of *Brookings Papers* includes the papers and discussions presented at the conference. The first paper presents and tests a new model of inflation that rejects the conventional theory of a natural rate of unemployment. It shows that a rate of inflation that is above zero but too low to factor into the decisionmaking of most workers and employers will permit unemployment to be sustained well below its so-called natural rate. The second paper examines "new economy" explanations for the recent spectacular rise in stock prices by comparing the ability of stock values and professional forecasts of U.S. firms' earnings to predict firms' investment behavior. The third paper applies growth accounting methods to recently revised official U.S. productivity data to analyze the sources of the recent surge in productivity and to offer informed judgments about whether it is likely to continue. The fourth paper investigates the possibility of a link between share prices, interpreted as an indicator of future profitability, and unemployment across a sample of industrial countries whose unemployment rates diverged during the 1990s.

SINCE THE 1970S many economists have accepted the idea of a natural rate of unemployment that describes a unique equilibrium for the real economy. The NAIRU, the empirical counterpart of the natural rate, has become part of the toolkit of policymakers and analysts, especially in the United States. And until the past few years, estimates of a relatively constant natural rate of about 6 percent have fit the actual behavior of the U.S. economy over the previous three decades reasonably well. Since then, however, the persistence of low inflation alongside unemployment rates that have fallen well below 6 percent has led most analysts to conclude that

the U.S. NAIRU has fallen. In the first paper of this volume, George Akerlof, William Dickens, and George Perry go further and reject the natural rate model itself.

Over the years there has been considerable controversy about how to model the formation of expectations in empirical work with the natural rate model. Akerlof, Dickens, and Perry, by contrast, see how people use expectations, rather than how they form them, as the key. Rather than accept the standard economic assumption that people make the best use of all available information, the authors turn to a variety of sources for evidence about how people actually use information in making decisions. Supported by this evidence, they develop an alternative to the natural rate model that is based on behavioral assumptions they find more realistic and that fit the facts better. A striking feature of their model is that it exhibits, rather than a unique natural rate, a range of sustainable unemployment rates consistent with steady, low rates of inflation. The authors show that the lowest unemployment rate in this range is well below the natural rate as usually estimated.

The authors cite psychological studies that have found that decision-makers "edit" the information available to them, ignoring much that is potentially relevant in order to concentrate on the few factors that matter most. Similarly, studies on the psychology of perception show that an event or stimulus must pass a threshold before it is even perceived, let alone acted upon. In addition, from interviews with compensation professionals, the authors infer that wage setters do not behave as most economic models assume. Rather than choosing a real wage target and then adjusting it fully for expected inflation, they mix information about inflation with a variety of other information relevant to wage setting. And from interviews of the lay public by Robert Shiller and questionnaire studies by Eldar Shafir, Peter Diamond, and Amos Tversky, they find telling evidence about how people perceive and react to inflation. Employees systematically underestimate the tendency of inflation to boost their own nominal wages. Therefore, the authors reason, in periods of moderate inflation, employees' job satisfaction is likely to be high even if their real wage is unchanged. Employees are pleased by their wage increase and do not fully recognize the corresponding rise in prices.

Drawing on these insights, the authors construct a model in which firms pay an efficiency wage and workers respond according to their view of job

and wage opportunities outside the firm. The crucial issue is how wage-setting behavior varies with the inflation rate. At any given time, some firms and their workers are fully rational, their wage setting fully incorporating expected inflation. At the same time, other firms or their workers are near rational, and their wage setting responds fully to current conditions in their labor market, but less than fully or not at all to expected inflation. Because all firms adjust to current market conditions each time they set wages, the wage level in near-rational firms trails the average wage, but only by a small, and not cumulating, percentage. Prices in all firms are a markup on expected unit labor costs. The authors show that the cost in lost profits from near-rational behavior is negligible at low rates of inflation and grows with the inflation rate. This supports their key hypothesis that the proportion of firms that fully adjust for expected inflation when they set wages will rise with the inflation rate. With little or no inflation, a large fraction of firms will be near rational, not fully adjusting their wages and prices in this way. At successively higher rates of inflation, that fraction will decline, and at a sufficiently high rate of inflation all firms will fully incorporate expected inflation.

The authors present a formal model that derives the implications of this firm-level behavior for the inflation-unemployment relation in the macroeconomy. At zero inflation, rational and near-rational behavior coincide, and equilibrium unemployment is at what would be the conventional natural rate. When steady-state inflation is between zero and some moderate rate, higher inflation is accompanied by lower unemployment. This trade-off reflects two effects working in opposite directions: as the inflation rate rises, firms pursuing near-rational behavior increase employment, but firms shifting to fully rational behavior cut their employment as they shift. In this inflation range, the first effect dominates. Beyond some inflation rate the second effect begins to dominate, and higher inflation is associated with higher unemployment, eventually approaching the conventional natural rate as all firms and workers fully incorporate inflationary expectations. There is thus a point of lowest sustainable unemployment. The natural unemployment rate, which is the rate associated both with high levels of inflation and with zero inflation, is significantly above the lowest sustainable unemployment rate. More generally, operating with an inflation rate either higher or lower than that associated with the lowest unemployment rate leads to excessively high unemployment in the long run. The short-run Phillips curve that follows from this model is

one in which the coefficient on expected inflation rises with the inflation rate and approaches unity at sufficiently high inflation rates.

The authors provide empirical support for this form of short-run Phillips curve from U.S. data over the period since the Korean War. They first discuss results reported elsewhere by William Brainard and Perry, who used general Kalman filter estimation that permits all Phillips curve coefficients to vary over time. These authors found little or no variation in the intercept or in the coefficient on unemployment. But the coefficient on lagged inflation, the conventional proxy for expected inflation, did vary over time. Starting from low values in the low-inflation years before the 1970s, it rose to its peak values with the high inflation that accompanied the two OPEC oil price shocks, and then declined again as the inflation rate fell to its recent low levels. Akerlof, Dickens, and Perry present their own results from fitting separate Phillips curves to periods of low and high inflation. To do this they sort the quarters from 1954 through 1999 into two samples: those quarters when the trailing five-year average of inflation as measured by the consumer price index was below 3 percent (or, alternatively, 2.5 percent), and those when it was above 4 percent. The samples have mean inflation rates of 2.0 percent and 6.3 percent, respectively. Using a variety of specifications, the authors then estimate their model with three alternative unemployment measures to allow for different treatment of demographic changes, three measures of price inflation, and a measure of wage inflation. They find the coefficient on inflationary expectations to be consistently and substantially larger in the high-inflation-period regressions than in the low-inflation-period regressions. The authors also report parallel regressions using direct survey measures of inflationary expectations, which avoid the ambiguity over whether lagged inflation is an adequate proxy for expected inflation. These estimates show an even sharper difference between the coefficients on price expectations in the low- and high-inflation periods. All the split-sample least-squares regressions thus support the key hypothesis of the formal model: how price expectations are incorporated into wage setting varies with the inflation rate.

The authors go on to derive an approximation to their model that is tractable for nonlinear estimation. In addition to unemployment and expected inflation, the conventional right-hand-side variables in a Phillips curve, the estimation model includes a term representing how past inflation affects the likelihood that people will act rationally toward expected infla-

tion. To provide a check on the robustness of the results, that term is proxied in a variety of ways. So is expected inflation itself, which is represented by alternative forms of distributed lags on past inflation as well as by the direct survey measures of expected inflation. In addition, the several alternative measures of unemployment and price and wage inflation used in the least-squares regressions are used in estimating the nonlinear model. Finally, whereas the full data period ran from 1954:1 through 1999:4, some of the nonlinear regressions were run only through 1989:4 to ensure that the results are not driven simply by the one long episode of low inflation and falling unemployment in the 1990s. The authors find that the estimates from the nonlinear regressions support the formal model.

The most important result is that nearly all the point estimates indicate that large, sustainable gains in employment are available by operating the macroeconomy at inflation rates moderately above zero. To summarize the range of estimates, for each of their regressions the authors calculate the inflation rate associated with maximum employment and the difference between the corresponding unemployment rate and the natural unemployment rate, which measures the employment gains available. Not surprisingly, the large number of specifications generates a wide range of point estimates. The densest cluster of estimates spans a range from 1.5 to 3 percent for the employment-maximizing inflation rate and from 0.5 to 3 percentage points for the corresponding unemployment reduction.

The authors discuss in detail four estimated equations that are representative of the range of their specifications, including one that is estimated only through the end of the 1980s. In each case the coefficient on inflationary expectations stays at or near unity starting some time after the onset of the inflationary period that began in the late 1960s, and takes on much smaller values during the low-inflation periods before and after. The amount of variation in this key parameter, as well as the timing of the variations, differs across the specifications, but all support the key prediction of the authors' model. The employment-maximizing rate of inflation for these equations ranges between 1.6 and 3.2 percent. And the reduction in unemployment below the natural rate ranges from 1.5 to 3.1 percentage points.

Although these estimates provide support for the qualitative features of the model, they do not pin down numerical targets for inflation and unemployment. Indeed, it would be unrealistic to seek precise estimates. However, their main results and the departures from conventional natural rate

models that they identify appear to be robust and do provide useful guides for econometricians and policymakers. Rather than a unique natural rate of unemployment, the economy has a range of sustainable unemployment rates that are consistent with moderate rates of inflation. For econometricians the main message is that models that assume a unique natural rate are misspecified and will yield misleading estimates of the economy's potential and of the minimum rate of unemployment that is sustainable. For policymakers the results provide some broad guides. Zero inflation is an inappropriate policy target, because it raises the sustainable rate of unemployment by an important amount. High rates of inflation are bad for the same reason, as well as because of the distortions and inequities they bring. Moderate rates of inflation, in a range that includes the experience of recent quarters when the core CPI has been rising at about a 2 percent annual rate, allow the economy to operate with low rates of unemployment and are consistent with a policy seeking to maximize prosperity.

THE 1990S WITNESSED unusually high rates of investment in plant and equipment by U.S. firms, yet even so the dramatic rise in stock market valuations during the period far outstripped all estimates of the growth in firms' stocks of tangible capital. Estimates of average q, the ratio of the market value of the average firm to the replacement cost of its tangible capital, already high by historical standards in 1990, roughly doubled during the decade. There seems little doubt that the extraordinary increase in household wealth that this rise in market value represents has been an important factor in the boom in consumption and the decline in saving from personal income that also characterized the 1990s.

But there is little agreement about the reasons for the rise in stock prices. Some view the 1990s as a period of irrational exuberance, of unrealistic expectations for growth in future earnings, or of a mistaken belief in the sustainability of capital gains that are in fact based on a speculative bubble. Others have greater faith in market efficiency, seeing in the rise the reflection of a "new economy," built on the power of information technology and its promise of future growth in productivity and profits. In this view, conventional measures of q are missing a recent accumulation of intangible capital—human or intellectual—that has become as important as the physical capital that was the source of profits in the old economy. Still others suggest that the rise reflects a reduction in risk premia, as investors have come to appreciate the ability of diversification across a

portfolio of stocks to reduce risks. These different views of the market's rise have drastically different implications for the future course of events. If the first view is correct, a day of reckoning will eventually arrive, and households will suffer dramatic reductions in their wealth. But if the second view is correct, we may reasonably expect this period of extraordinary prosperity and growth to continue. And if the third view is the right one, we may have reached a new wealth plateau, with no reason for euphoria looking ahead but also no reason to expect a sharp decline in wealth; rather we should expect *q*s to gradually return to normal as the reduced risk premia stimulate investment and capital formation. In the second paper of this issue, Stephen Bond and Jason Cummins attempt to assess the plausibility of these different views of the growth in market values, using data on investment, market valuations, and professional earnings forecasts for a sample of more than 1,100 firms.

The authors' strategy is to determine whether firms treat market values as a source of information relevant to their tangible investment decisions, taking account of the possibility that these values may in part be signaling the profitability of intangible investment. A major ingredient in the authors' analysis consists of estimates of the present value of future earnings constructed from earnings forecasts by professional securities analysts. The earnings forecasts are supplied by I/B/E/S International, a private company that has been collecting such forecasts since 1971. The present values calculated from these forecasts, the authors argue, are potentially a better measure of the future profitability of a firm than are market valuations. For a given year, the authors average the analysts' one- and two-year-ahead earnings forecasts for each firm and then grow this average at the analysts' long-term growth forecast rate for the firm for the subsequent three years. Beyond five years the firm's earnings are assumed to grow at the economy's average growth rate over the entire sample history. To obtain present values, these nominal earnings forecasts are discounted by the current nominal rate on long-term Treasury bonds plus a fixed risk premium of 8 percent for the first five years. For consistency with the assumption that the longer-term earnings of individual firms return to the economy's sample average, the discount rate beyond five years is based on the mean interest rate for the sample period. Dividing the resulting present value by the replacement cost of the firm's capital provides the authors with a q based on analysts' earnings estimates, which they call $\hat{q}$.

Although $\hat{q}$, like the estimate of q based on equity values, q^E, has increased dramatically, there is a significant difference in the behavior of the two measures. Whereas the average value of $\hat{q}$ grew approximately 200 percent over the period from 1982 to 1998, average q^E grew 330 percent over the same period. Although the average values moved in tandem in a number of years, their changes are only loosely correlated, with a correlation coefficient of 0.15. These differences are mimicked at the firm level: median q^E was about 15 percent above median $\hat{q}$ in 1982, whereas median q^E was roughly 75 percent above median $\hat{q}$ by 1998. The authors display a plot and a nonparametric regression of $\hat{q}$ against q^E. Although the two measures are positively related, neither the correlation nor the slope is near unity, and in the range where most of the observations are located, the expected value of $\hat{q}$ is almost constant at one. Differencing the two variables, which should diminish the importance of slowly changing measurement errors in general, or errors in the authors' discounting procedure in particular, yields qualitatively similar results. Some features of the data do stand out. It appears that the market's assessment of companies has become more heterogeneous than those companies' earnings would warrant and that both the mean and the median have increased much more for q^E than for $\hat{q}$. Other noteworthy features of the data are that the median level of intangible investment is zero, and that, although intangible investment is much more variable across firms than tangible investment, the rate of intangible investment does not vary much among those firms that engage in it.

The authors believe it is hard to argue that investors at large have a more accurate view of future earnings than the professionals who study the firms, and they therefore regard the enormous growth in q^E relative to $\hat{q}$ as supporting the view that the market is overvaluing assets. But it is possible that earnings forecasts underestimate the future returns to intangibles. In any case, the denominator of both qs should adjust for the growth in intangible assets, which are omitted in the denominator of both measures. The difficulty of proving or disproving this hypothesis, of course, is that no direct measure of the stock of intangibles is available. The authors therefore propose an indirect test, taking advantage of the theoretical connection between market valuation and investment. According to standard q investment theory, the higher is q, the more firms will invest. They note that this theory has not fared particularly well, however: coefficients of q in investment equations tend to be small, implying implausibly high

adjustment costs, and imprecise. A variety of explanations have been offered for these failings, including capital market imperfections and nonconvex adjustment costs. The authors suggest two other possibilities that would give spurious results: that the stock market gives noisy signals of firms' fundamentals, and that intangibles are important contributors to market values. If a firm's managers perceive the market as giving noisy estimates of the firm's profit opportunities, their response to changes in market valuations will be less than implied by the theory. And even if firms believe the market gives accurate estimates of their valuations, there is no reason for a firm's investment in tangibles to respond to changes in the value of its intangibles.

The authors' estimates of $\hat{q}$, already discussed, provide one of the ingredients needed to examine the possibility that noise in stock market valuations is a culprit. But to deal with the second possibility the authors need to extend conventional q theory to deal with the possibility that two different types of capital make important contributions to the firm's value, only one of which is measured. Their model, following the approach of Fumio Hayashi, assumes that the profit function is linearly homogeneous in the two capital stocks and the rates of investment, and that their costs of adjustment are additively separable and quadratic. Under these assumptions, investment in tangible capital depends linearly on the ratio of the firm's total market value to the replacement value of its tangible capital, on the ratio of intangible investment to tangible capital, and on the ratios of the stocks and values of the two types of capital. If it is assumed that the ratios of the two types of capital and of their prices are relatively stable for a given firm, those terms in the expression can be replaced by a firm-specific fixed effect. What remains is familiar: tangible investment is a positive function of q, with a coefficient inversely proportional to tangible capital's adjustment cost, but with an additional term that subtracts the ratio of intangible investment to tangible capital. This term's coefficient reflects relative costs of adjustment of the two types of capital and their depreciation rates. The intuition behind subtracting intangible investment is straightforward. A firm making intangible investment should have a high value of intangible capital. Subtracting the weighted value of intangible investment "adjusts" the market value for the value of intangible capital, leaving the valuation of tangible capital appropriate for explaining tangible investment. Although this version of the investment equation avoids the need for a direct measure of the stock of intangibles, it does require a

measure of investment in intangibles. In their empirical implementation, the authors use research and development (R&D) and advertising as variables that are plausibly related to tangible investment, even though they are expensed on the firm's income statements. Of course, the variables on the right-hand side of the equation are endogenous and may well be correlated with the error in the equation; the authors deal with this difficulty by using instrumental variables in their estimation.

Rather than use the simple q^E measure, which reflects only the value of equity in a firm's investment equation, the authors calculate the theoretically more appropriate measure, denoted Q^E, which takes into account the firm's debt and the effect of taxes and depreciation allowances. To calculate $\hat{Q}$, the comparable measure using professional analysts' earnings forecasts, they simply replace the stock valuation with the present discounted value of earnings previously estimated.

The authors estimate investment equations for tangible investment, using various combinations from the menu of explanatory variables Q^E,$\hat{Q}$, R&D, and advertising. The equations are estimated by the generalized method of moments (GMM) in differenced form and subjected to specification tests. To deal with the issue of endogeneity, three- and four-period-lagged values of the dependent variable and cash flow are used as instruments. Consistent with earlier work, the equations using Q^E do not fare well. Included by itself, Q^E has a small and marginally significant coefficient, and the equation fails the specification test. When advertising and R&D are added as variables, they have the right sign and are significant (either singly or together), but they only moderately improve the performance of Q^E.

Substitution of $\hat{Q}$ for Q^E gives dramatically better results. Coefficients on $\hat{Q}$ are five times those on Q^E and are highly significant. Advertising and R&D both have the appropriately negative coefficient when entered separately, but only advertising is statistically significant; when all variables are included, advertising retains its importance but R&D's coefficient becomes essentially zero. This version of the equation also passes the specification tests. Strikingly, when both $\hat{Q}$ and Q^E are included, $\hat{Q}$ is found to do all the work: market valuations add essentially no information to that contained in the analyst-based measure of a firm's value. This is true whether or not the variables proxying for intangible investment are included. However, the much stronger performance of R&D in the Q^E equations is consistent with the possibility that professional forecasters do not anticipate the returns to R&D as well as the market does.

To check on the robustness of the results suggesting that share prices are noisy indicators of the profitability of firms, the authors undertake three other analyses of the data. First they focus on a sample of firms that are intensive in intangible investment, defined as those firms whose intangible investment (scaled by their tangible capital) is in the top quartile of all firms. For all subsamples of this set of firms, Q^E does poorly but $\hat{Q}$ has a substantial and significant coefficient. For the subsample of firms whose advertising expenditure is in the top quartile, the coefficient for advertising remains significant. However, the coefficient is only about half as large as that for the entire sample, suggesting the possibility that differences in the relationship of tangible investment to Q for advertising-intensive and advertising-nonintensive firms influences the coefficient for the full sample. For the subsample of firms with R&D expenditure in the top quartile, on the other hand, the results are more supportive of the view that intangible investment is an important source of the story; for these firms the coefficient is substantial and significant. As a second check on robustness, the authors divide the sample into three groups of roughly equal size on the basis of the proportional deviation of Q^E and $\hat{Q}$. Although Q^E by itself does somewhat better for the sample with the smallest deviation than for the entire sample, inclusion of $\hat{Q}$ continues to obliterate its influence.

The authors perform a third check on whether the differences between market valuations and estimates of the present discounted value of earnings are mostly noise or reflect intangible capital and a systematic underestimation of the earnings from such capital on the part of professional analysts. The authors examine the data for "new economy" companies, defined as those engaged in the manufacture of computers or other electronics products or in software or telecommunications, and "old economy" companies; they also examine the data by manufacturing industry. They find that the enormous dispersion in the difference between market and analyst-based qs is pervasive, appearing within all industries, as is the growing gap on average. Interestingly, if anything, investment tracks q^E better for new economy firms than for old. But as in their earlier analysis, it is clear that $\hat{q}$, based on earnings forecasts, predicts tangible investment far better for both.

The authors believe that the wide and growing gap between market valuations of firms and valuations based on expected future profits is more likely to reflect noise in stock prices than systematic errors by professional forecasters. Analyst-based valuations are demonstrably more informative

about firms' tangible investment behavior than market valuations, and although the authors identify a limited role for intangible investment, they do not believe it can account for the spectacular rise in the stock market valuation of firms. They regard as their most surprising finding the fact that there appears to be no information about investment behavior in market valuations once expected future earnings are taken into account. Managers appear to make investment decisions on the basis of fundamentals and do not respond to movements in share prices that do not reflect these fundamentals. In this regard, their results do not support simple *q* theories of investment. The authors conclude that although policymakers may be right to worry about the wealth effects on consumption from a reversal of the recent dramatic appreciation of market values, they need not be unduly concerned about the impact of "irrational exuberance" on business investment.

THE RECORD PEACETIME EXPANSION in the United States has been accompanied by a fall in the unemployment rate to its lowest level in thirty years, a return to virtual price stability, and an unprecedented appreciation of the market valuation of firms. During the past five years of this expansion, when productivity growth might have been expected to slow, it has instead quickened to rates almost matching the post–World War II peaks of the 1960s. This extraordinary growth in productivity itself helps explain the decline in inflation and the rise in stock market values. As the expansion continues to break records, a consensus has emerged that something fundamental has changed, but not on the causes of that change or the likelihood of its permanence. "New economy" proponents credit the success to the information revolution, which they see as driving a fundamental transformation of the economy that will lead to faster productivity growth for many years. Skeptics of this view acknowledge the importance of the high rates of investment stimulated by the computer revolution, but attribute much of the economy's success to a series of favorable but temporary shocks. In the third paper of this issue, Dale Jorgenson and Kevin Stiroh analyze information made available by the recent benchmark revision of the U.S. national income and product accounts to examine the recent spurt in productivity. Their results, based on a standard growth accounting framework, allow them to make an informed judgment about the likelihood that the extraordinary recent surge in productivity growth will continue.

The authors report that, after a twenty-year slowdown, growth in average labor productivity (ALP) during 1995–98 was 2.4 percent per year, roughly a percentage point faster than over the preceding twenty years. Although ALP is of intrinsic interest, relating directly to the standard of living, its growth reflects both growth in nonlabor inputs and improvements in technology. Since business investment during the past half decade has been unusually strong, much of the acceleration of ALP could have come from capital deepening. To try to disentangle the sources of growth, the authors use the framework based on the production possibilities frontier developed earlier by Jorgenson. This framework differs from one using a Solow-style aggregate production function in that it recognizes multiple outputs and can capture substitutions among outputs of investment and consumption goods, as well as among inputs of capital and labor. The authors argue that a framework like theirs, which allows the relative prices of outputs to change, is essential for this endeavor, given the significant decline in the prices of computers and software relative to the prices of other investment goods. Assuming competitive product markets and constant returns to scale, growth in total factor productivity (TFP), analogous to the Solow growth residual, is measured as the difference between the sum of the share-weighted growth of outputs and the sum of the share-weighted growth of inputs.

The authors' aggregate output concept is similar to that used by the Bureau of Labor Statistics, but differs in that it includes imputations for the service flow from consumer durables and owner-occupied housing. On the output side, the authors separately estimate computer and software purchases by consumers (which they label consumption), and on the business side, consumer services from computers and software, computer investment, software investment, and communications investment. On the input side, in addition to labor, they separately calculate the contributions of the business services from computers, software, and communications capital, and the consumer services from computers and software and from owner-occupied housing. The measures of stocks and flows of capital services are computed from detailed price and quantity information on fifty-two types of nonresidential assets, five types of residential assets, and thirteen types of consumer durables.

The most striking feature of the data is the very rapid relative price decline for computer investment: after averaging 18 percent per year from 1960 to 1995, this price decline accelerated to over 27 percent per year.

The relative prices of software and communications equipment, the other categories of information technology (IT) capital, are relatively flat early in the period but begin to fall in the late 1980s. Investment in computers and software capital, which was nil as recently as 1960, has now overtaken investment in communications capital. Although business investment in IT capital is far larger than IT investment by households, households have spent substantial sums, more than $20 billion per year, on computers and software since 1995. Yet even with this high rate of IT investment, the stock of IT capital remains a relatively small part of the economy's total capital stock. Using a broad definition of capital, including tangible assets such as equipment and structures, consumer durables, land, and inventories, the authors estimate that the entire stock of capital in 1997 was worth $26 trillion. In 1998, IT assets amounted to only 3.4 percent of that tangible capital stock and 4.6 percent of reproducible tangible capital.

In the authors' methodology, output depends on the flows of capital services, which differ per unit of stock for different types of capital. Capital services of a given type of capital are assumed to be equated by firms (or households in the case of consumer durables) to the capital's rental price. That price reflects the prevailing interest rate, depreciation (including obsolescence), and gains or losses on capital due to changes in its price, all adjusted for taxes. Using rental prices for each type of capital, the authors calculate the flows of services for the total capital stock and for each of the high-technology asset classes. For 1998 the authors estimate that capital services were only 12.4 percent of the capital stock for tangible assets as a whole, but 40 percent of the stock of IT capital, reflecting the rapid price declines and high depreciation rates that affect the rental prices of IT capital. Their calculations show a clear increase in the growth of aggregate capital services, from 2.8 percent per year for 1990–95 to 4.8 percent for 1995–98. Although the stock of IT assets is relatively small, they are estimated to have provided 11.2 percent of total capital services in 1998, and the acceleration in total capital services largely reflects their growth. In that year the share of services from computer hardware alone reached 3.5 percent of the total. According to the official price indexes, software had smaller price declines than computer hardware and therefore had lower rental prices. Nevertheless, because of high rates of investment, real capital services from software are estimated to have grown at 13.1 percent per year in the 1990s.

To complete their growth accounting, the authors need estimates of labor input from 1959 to 1998. They distinguish labor input from hours worked, taking account of changes in the composition of hours worked as firms substitute among heterogeneous types of labor. To do so, the authors construct constant-quality indexes for labor input and its price, taking into account the sex, employment class, age, and education level of workers. According to their estimates, growth in labor input increased to 2.8 percent per year over 1995–98 from 2 percent over 1990–95, primarily because of growth in hours worked, as labor force participation increased and unemployment rates fell. This growth in hours was partly offset by a decline in the growth of labor quality of about 0.2 percentage point, reflecting demographic trends and exhaustion of the pool of available workers.

Armed with their estimates of the service flows of labor and the various types of capital and their shares, the authors calculate the contribution of capital and labor services and of TFP to U.S. output growth. Output as measured by the authors has grown at an average annual rate of 3.6 percent over the past forty years. According to their decomposition, growth in inputs accounts for about 80 percent of this growth: capital services contribute an average of 1.8 percentage points, and labor services 1.2 percent. TFP improvements account for the remaining 0.6 percentage point. Capital accumulation and growth in hours accounted for approximately three-quarters of the contribution of capital and labor, respectively; growth in the quality of capital and labor was also important. After the rapid growth of output and TFP in the 1960s and early 1970s, both slowed markedly through 1990, with TFP growth falling almost two-thirds of a percentage point, from 1 percent to 0.3 percent. Growth in capital inputs also slowed by more than a percentage point, further contributing to the decline in ALP from 2.9 percent for the period 1959–73 to 1.4 percent for 1973–90.

The turnaround during the 1990s was quite remarkable. Relative to the early 1990s, output growth increased by nearly 2 percentage points in 1995–98, with capital's contribution jumping by 1 percentage point, labor adding another 0.4 percentage point, and TFP quickening by 0.6 percentage point. The authors document the rising contributions of IT to growth. Their calculations show that the recent surge in IT investment and consumption nearly doubled the output contribution of IT for 1995–98 relative to 1990–95. IT now accounts for more than 40 percent of the total growth

contribution from broadly defined capital. The authors regard the improvement in TFP growth as perhaps the most remarkable feature of the data, suggesting massive improvements in technology and increases in the efficiency of production.

The rapid capital deepening and TFP growth in recent years translate into striking growth in ALP in the late 1990s. The authors' estimates imply that more-rapid capital deepening added 0.49 percentage point to ALP growth, with the improvement in TFP adding another 0.63 percentage point. As previously mentioned, a decline in the growth of labor quality partially offset these improvements.

The authors find it plausible that official price indexes for software and telecommunications equipment overstate the increases (or understate the declines) in the prices of those goods. They examine the implications for growth of two alternative price scenarios, one with a moderate price decline and one with a rapid price decline. The moderate case assumes that the rate of decline of prepackaged software prices, which are adjusted for quality in the official indexes, occurred in all other types of private software investment, including custom and business own-account software. The moderate case also assumes that the prices of digital switching equipment fell at a 10.7 percent annual rate, a number suggested by earlier work by Bruce Grimm. The rapid price decline case, based on estimates by other authors, assumes that prices for software and communications equipment fell 16 percent per year in the period 1959–98. The results for the two cases are not surprising: relative to their base case, output growth for 1995–98 is somewhat greater (0.16 and 0.34 percentage point for the moderate and the rapid price decline cases, respectively), the contribution of software services is greater by similar amounts, and the TFP residual grows slightly less.

What sectors of the economy are the sources of TFP growth? There is no doubt that increased productivity in the computer industry itself has made important contributions to aggregate growth, with estimates by the Council of Economic Advisers and others suggesting a contribution to TFP of 0.4 to 0.5 percentage point for the period 1995–99. The authors' own calculations give similar numbers for the base case, but they stress that the results are sensitive to assumptions about the price declines in IT. Under the moderate decline case assumptions, the contribution of IT for this period increases to 0.64 percentage point, and in the rapid decline case the contribution of the IT sectors to total TFP growth for 1995–98 is

roughly doubled, to 0.86 percentage point. But as previously indicated, overall increases in TFP growth were smaller because TFP growth in other sectors was smaller, reflecting the increase in services of IT capital. The authors' analysis, in contrast to recent work by Robert Gordon, finds that the increases in TFP during the late 1990s are not entirely located within IT-producing industries. Non-IT growth increased markedly in each of their three cases. The authors note that this could be interpreted as evidence of a "new economy," with spillovers from IT into non-IT industries. But it could also simply reflect technological progress in non-IT industries that is entirely independent of the IT revolution.

To trace aggregate TFP growth to its sources in the productivity growth of individual industries, and to measure the effects of reallocations of outputs and inputs among industries, the authors turn to detailed data on thirty-seven industries for the period 1958–96, compiled from a variety of sources. They note that their industry data are not strictly comparable to the aggregate data analyzed above; most notably, they do not reflect the benchmark revisions published by the Bureau of Economic Analysis in 1999, and thus show slower output and productivity growth. For each industry the authors estimate the contribution of inputs and productivity growth using a gross output concept that includes purchases of intermediate inputs from other industries as well as the primary inputs from capital and labor services of a value-added concept. Since, for the typical industry, gross output considerably exceeds value added, the sum of gross output across industries exceeds the sum of value added, and the contribution of a particular industry's gross output productivity improvements to economy-wide TFP growth is magnified as it works its way through the production process. Consistency in measuring these improvements requires using so-called Domar weights (weights that add up to more than one) to aggregate industry-level productivity to aggregate TFP growth.

The authors find that industry productivity growth was highest in two high-technology industries: industrial machinery and equipment, and electronic and electric equipment. The first category includes the production of computers, and the second the production of semiconductors and communications equipment. This finding fits the story of enormous technological progress in these high-technology capital goods, which has generated falling prices and substitution toward IT in other sectors. Indeed, correctly taking into account the declining price of semiconductors, which are produced by the electronic equipment industry, and the increased flow

of these intermediate goods to the industrial machinery industry lowers the productivity increase credited to the latter and rightly allocates the progress to the former. These two industries also show dramatic increases in labor productivity, with rapid accumulation of capital and growth in intermediate inputs. The authors' calculations also show relatively strong productivity growth in agriculture, textiles, rubber and plastic, instruments, and trade, but, somewhat perplexingly, negative growth rates in nine industries. Some other authors have reported similar findings and suggested that they reflect persistent measurement problems, but Jorgenson and Stiroh offer reasons why they may be actual declines. They therefore suspend judgment until more careful research linking firm- and plant-level productivity to industry productivity has been done. They also comment that several of the industries with slow productivity growth are heavy investors in IT, a fact that should disconcert "new economy" proponents, who argue that the use of IT is fundamentally changing business practices and raising productivity throughout the economy.

What is the likely prognosis for U.S. economic growth? It is clear that falling unemployment and higher labor force participation cannot fuel growth indefinitely. Indeed, many observers would agree that current high employment rates are unsustainable. But the recent improvement in TFP may be more likely to persist, and if prices of IT capital continue their rapid decline, we may continue to reap the benefits of rapid capital deepening. The authors organize their discussion of future prospects around the growth projections of the Congressional Budget Office (CBO), which are based on a relatively sophisticated multisector growth model of the U.S. economy. The CBO projects growth in potential GDP at 3.1 percent per year for 1999–2010, and actual growth at 2.8 percent, under the assumption that the economy will move to somewhat lower levels of utilization of labor and capital. For the nonfarm business sector, which is the most common basis for productivity analysis, the CBO projects potential growth of 3.5 percent a year.

Although they accept the CBO's relatively optimistic assessment of continued technological progress that underlies this projection, Jorgenson and Stiroh suggest some adjustments to the building blocks of this potential output projection. They note that the CBO does not adjust labor input for changes in labor quality, so that composition changes are included in their TFP projections and essentially held constant. They also question the CBO's projected 4.4 percent growth rate in capital inputs because it

exceeds the CBO's projection of nonfarm output growth by much more than that seen even in recent experience. Making their own adjustments to the contributions from capital, labor, and productivity, the authors project potential annual nonfarm output growth at 3.35 percent, only modestly below the CBO's projection. They stress, however, that any attenuation of the price declines that have resulted from rapid productivity growth in the IT-producing sectors would reduce these projections of output growth through two channels. First, TFP growth in the IT-producing industries themselves would slow, and second, investment by IT-using industries would slow, reducing the rate of improvement in labor productivity that has come from capital deepening.

FOR MOST OF THE first twenty-five years of the postwar period, the rich economies of the world prospered together. When recessions interrupted expansions, they were relatively brief, and there was little if any upward trend in unemployment rates. At the end of 1973 the U.S. unemployment rate was near 5 percent, and unemployment rates in the major countries of Europe averaged about 3 percent. The bad times that followed also hit all the OECD economies together: each experienced a deep recession following the first oil price shock and another following the second oil shock at the end of the 1970s. Since then, however, economic performance has diverged widely, with double-digit unemployment rates persisting in some countries for many years while unemployment fell in others. In the fourth paper of this volume, Jean-Paul Fitoussi, David Jestaz, Edmund Phelps, and Gylfi Zoega exploit this diversity of recent experience to try to understand why some advanced economies have performed so much better than others and, in particular, why unemployment persisted at historically high rates in much of Europe throughout the 1990s.

Several explanations have been offered for the high unemployment rates experienced throughout Europe in the 1980s. Some focus on restrictive labor market institutions and the work disincentives inherent in social safety nets. Others cite an excessive rise in real wages that resulted from slowing productivity and strong labor unions. Phelps himself previously modeled rising unemployment as a consequence of an increase in the effective cost of capital—the difference between real interest rates and expected productivity growth—that came from expectations of slower productivity growth and the relative rise in nonwork income available from social programs and rising wealth. The present authors argue that devel-

opments in the 1990s cry out for a reexamination of these ideas, since the proposed explanatory variables have changed in some countries and performance has improved in some, but the correlations are not always the predicted ones. They also note that several new hypotheses have been offered, including the importance of labor market reforms and monetary policies, and that some important new shocks have occurred, including rising stock markets and revivals of productivity growth in some countries. The authors believe these new shocks are interrelated and are consistent with their hypothesis that "new economy" effects have been important in some countries but not others, and that this helps explain the divergent unemployment experience across countries. An important feature of the authors' approach is that the NAIRU for an economy is modeled as an endogenous variable. The NAIRU thus represents a moving equilibrium path that changes relatively freely in response to various shocks, rather than a feature of an economy that responds only to significant structural changes, and then only gradually.

The authors conduct their main empirical examination of these ideas and developments within the framework of Phelps's earlier formal model, which focused on two main variables in explaining unemployment: the effective cost of capital and income from private wealth. In an equation estimated across decade averages, these two variables taken together do a reasonable job of tracking the change in average unemployment rates by country between the 1970s and the 1980s and between the 1980s and the 1990s. The authors follow these preliminary results with more detailed econometric investigations that constitute the main part of their paper.

The authors estimate a baseline specification rooted in their formal model and then consider variables suggested by other hypotheses that may interact with the variables in that model to better explain the changes in unemployment across countries in the 1990s. They first estimate separate regressions explaining each country's unemployment rate using annual data for the period 1960–98 for nineteen OECD countries. The main explanatory variables of interest are smoothed productivity growth, the average long-term real interest rate in the Group of Seven countries (as a proxy for world real interest rates), and the ratio of after-tax income from private wealth to productivity. They also include variables for lagged unemployment, the real price of oil, and the change in the inflation rate, which they treat as a proxy for demand effects. The estimated effects of the main variables vary considerably across countries. The authors hypothe-

size that this variation largely reflects differences in the degree of real wage rigidity, which affects the intensity with which changes in the main variables affect unemployment. Therefore, in their preferred formulation, they use a pooled regression in which they constrain the main variables to have the same effect in each country up to a factor of proportionality, which they call the country sensitivity coefficient. Regressions with this constraint and fixed effects produce common coefficient estimates on the main variables that have the expected sign. Although the interpretation of these results depends crucially on the constraint, they are within the ballpark of estimates that others have reported, and they appear plausible to the authors.

In these constrained estimates, differences across countries are captured in the fixed country effects and three parameters: the sensitivity coefficients, the persistence effects (measured by the coefficient on lagged unemployment), and the coefficient on the change in the inflation rate. Although all parameters have the expected sign, their magnitudes range widely across the nineteen countries in the sample, so that a shock of given size would be predicted to have widely different effects on unemployment in different countries. The authors illustrate these differences by comparing the predicted effects of changes in world real interest rates and country productivity trends across five major countries. According to their estimates, in the long run a 5-percentage-point rise in real interest rates would raise the unemployment rate by 0.5 percentage point in the United States and 1.8 percentage points in France. A 3-percentage-point decline in productivity growth would raise the U.S. unemployment rate by 1.3 percentage points and the French rate by 4.8 percentage points. Effects for Germany, Italy, and the United Kingdom are between these two in both cases, but nearer the estimates for France. The immediate effects of these shocks are much smaller but vary across countries in much the same pattern as these steady-state effects.

The authors then perform a direct check of the ability of their specification to explain unemployment changes: they use estimates made over 1960–91, a period during which unemployment rates diverged widely across the nineteen countries, to predict actual unemployment in 1998. The results are mixed. Actual unemployment in 1998 is within 1.5 percentage points of the predicted level in six countries, more than 1.5 percentage points higher than predicted in eight, and more than 1.5 percentage points lower than predicted in five. The authors conclude that, although their

formal model offers a reasonable description of how their main variables affect unemployment over much of the period, developments not captured by the model are also important in explaining the evolution of unemployment rates in the 1990s. They then turn to an examination of hypotheses that include some of these other developments.

Labor market reforms, which the OECD has championed, are one widely cited candidate for explaining the decline in unemployment in those countries that have pursued such reform. The authors show that variables measuring certain characteristics of the social safety net and labor markets help explain the fixed country and sensitivity effects estimated with their model. This finding suggests that these variables have a role in determining the average rate of unemployment in a country and how shocks have an impact on that rate. The kinds of variables that seem to matter are the replacement ratio and the duration of unemployment benefits, the proportion of the work force unionized or to which union wages apply, union coordination, active labor market expenditure by governments, and rules protecting employment. However, although these variables help explain differences in unemployment rates across countries in the 1980s, changes in these measures are much less effective in explaining changes in unemployment between the 1980s and the 1990s. When entered along with the authors' baseline variables, only union density and union coordination remain significant, and the authors report that the significance of coordination arises solely from the inclusion of Finland, where coordination fell while unemployment rose.

The authors next briefly explore possible monetary explanations for unemployment. To do this they use a variety of inflation and interest rate variables to explain the cross-country variation in the change in unemployment between the 1980s and the 1990s. They find that a decline in inflation is associated with an increase in unemployment but explains only a small part of the observed variation. High real interest rates in a country are associated with high unemployment and explain somewhat more of the variation than inflation does. But other indicators are less successful. They also find no effect on unemployment from staying in the European monetary system during the 1990s, which for some countries would have meant an externally imposed tight monetary policy. Taking all their findings into account, the authors infer that monetary policy may have had some effect on the evolution of unemployment. They suggest that the runup to European monetary union and the restrictive policies of the Bundesbank in the

aftermath of German unification may have caused unemployment to exceed its natural path during much of the 1990s, although they regard the evidence as inconclusive. They also note that the tight-money period appears to be over, with the high-unemployment countries no longer experiencing high real interest rates relative to their low-unemployment neighbors.

The authors' preferred explanation for diverging unemployment rates in the 1990s is that rising prospects for profits with the emergence of the "new economy" have accelerated job creation in some countries. They imbed that hypothesis in their baseline model by including stock prices, adjusted for the overall price level, as a measure of expected productivity and profitability. A crucial and distinctive feature of their formulation is that the expectation of improved productivity in the future is sufficient to increase employment and real wages in the present. Because trained labor is treated as a firm-specific asset, higher productivity in the future raises the present value of that labor in their model and thus expands firms' demand for labor. Additionally, the higher present value of other assets of the firm encourages investment in those assets and thus may expand the demand for the labor that works with them. As the authors put it, "In our model, firms decide to hire new workers when they become more optimistic about future profitability, even when they only want to maintain their current level of output."

To conform to the formal model's specification of labor demand and supply, in their empirical work the authors divide the real stock price by productivity, creating a ratio that they refer to as the normalized share price. This variable serves as a proxy for both the effective cost of capital and the profitability of an employee working with his or her equilibrium stock of the firm's other assets. When the normalized share price for each country is added to the earlier baseline regressions explaining annual unemployment rates over 1960–98, share prices have the expected negative sign and are significant, as are the other main variables of the model. Some variations on the normalized share price variable, such as the country share price relative to the OECD average, give similar results. Looking at decade averages, the authors find share prices to be significant when added to the earlier cross-country equations explaining the change in unemployment between the 1980s and the 1990s. They also find a rank correlation of –0.6 between the change in share prices and the change in unemployment rates between the 1970s and the 1990s.

Turning to a few of the individual countries in their sample, the authors leave their formal model to investigate the simple correlation between normalized share prices and unemployment over five-year intervals starting in 1960. They show graphically that a close and contemporaneous relation exists for the United States. Some relation is also apparent for Ireland, the Netherlands, and New Zealand, although it is less close, and five-year lags are often apparent whereas the authors' model predicts an immediate response of employment. And for four countries in which unemployment in the 1990s was higher than predicted by the annual baseline regression described above—France, Germany, Italy, and Spain—there appears to be a relation in only one (Spain), and it is long-lagged. The relation is unclear in the other countries and has the opposite of the expected sign over the last fifteen years in France. The authors see these mixed results as lending some support to their "new economy" hypothesis.

Comments at the Brookings Panel conference stressed the difficulty of identifying any causal role for stock prices, since these are presumably moved by any shock, including demand shocks that expand employment, as well as by changes in investor sentiment that may or may not prove warranted. However, if the authors are right and the strikingly improved performance of European stock markets in 1999 portends a surge in future profitability, the implications are important. We should then expect a speedup in growth and substantial reductions in unemployment in the major economies of continental Europe, driven by a speedup in productivity growth like that the U.S. economy has experienced.