Editors' Summary

THE BROOKINGS PANEL ON Economic Activity held its seventy-fourth conference in Washington, D.C., on September 5 and 6, 2002. This issue of *Brookings Papers on Economic Activity* includes the papers and discussions presented at the conference. The first paper reviews the process and methods of inflation and output forecasting at four central banks and proposes strategies for improving the usefulness of their formal economic models for policymaking. The second paper analyzes the implications for monetary policymaking of uncertainty about the levels of the natural rates of unemployment and interest. The third paper examines reasons for the recent rise in current account deficits in the lower-income countries of Europe and the role of economic integration in breaking the link between domestic saving and domestic investment. The fourth paper applies a new decomposition of productivity growth to a new database of income-side output to examine the recent speedup in U.S. productivity growth and the contribution made by new economy industries.

DESPITE A GROWING TRANSPARENCY IN the conduct of monetary policy at many central banks, little is still known publicly about the actual process of central bank decisionmaking. In the first paper of this issue, Christopher Sims examines this process, drawing on interviews with staff and policy committee members from the Swedish Riksbank, the European Central Bank, the Bank of England, and the U.S. Federal Reserve. Central bank policy actions are inevitably based on forecasts of inflation and output. Sims' interviews focused on how those forecasts are made, how uncertainty is dealt with, and what role formal economic models play in the process. In the case of the Federal Reserve, a history of subjective and model-based forecasts is publicly available, allowing Sims to evaluate and compare their performance both with each other and with private forecasts. He offers observations about how the performance of large econo-

metric models could be improved, and how statistical inferences from them could be made more useful to policymakers.

Sims begins with an informative overview of the process that takes place at each of the central banks before policy decisions are made. Within each institution, a policy round proceeds through a number of meetings among staff through which a forecast is arrived at iteratively, but the number of meetings, the way discussions are ordered, and the involvement of policy board members all vary. At the Federal Reserve, staff prepare the forecasts without policy board participation. At the other central banks there is typically some involvement of the policy board, and at the Bank of England some board members attend the meetings from the earliest stages. Each central bank has a primary macroeconomic model (the Riksbank has two) that is used to construct projections of the economy conditional upon various assumptions about future disturbances or policies; each also informally uses one or more secondary models. Sectoral experts play a major role at all four central banks. At the Federal Reserve their forecasts are compared with those from the primary model, and discrepancies serve as a basis for further discussion and possible adjustment of the primary model's forecast. The Federal Reserve's Green Book and the Bank of England's inflation reports summarize the results of the forecasting process and present a variety of potential time paths for the economy, which serve to communicate forecasting uncertainty. Other central banks also typically consider a range of outcomes.

Sims examines the situation at the Federal Reserve in greater detail than the others, analyzing the accuracy of the staff forecasts for inflation and output. These are publicly available, with access restricted only by a five-year nondisclosure rule. He expands in several ways on a recent study of the judgmental Green Book forecasts by Christine Romer and David Romer, adding four or five years of data, including the Federal Reserve's model-based forecasts in the assessment, and comparing both sets of forecasts with private forecasts. Through 1995 the model-based forecasts were based on the MIT-Penn-SSRC (MPS) model, developed in a collaborative academic effort and maintained by the Federal Reserve staff. Since then, model forecasts have been based on the new FRB/US model developed within the Federal Reserve. Sims compares both the Green Book and the model-based forecasts with forecasts from the Survey of Professional Forecasters (SPF) and with a naïve forecast that assumes no change from the current quarter.

Forecast accuracy can be characterized in different ways. The most straightforward is to calculate the root mean square errors (RMSE) of forecasts from actual outcomes, which Sims does for each of the four quarters following the forecast. For the inflation forecasts, he finds, as did the Romers, that the Green Book has the smallest errors at all horizons. The MPS forecast consistently finishes in second place except in the case of the one-quarter forecast; the naïve forecast does better than the SPF at every horizon. (The naïve forecast has the advantage of utilizing actual data for the current quarter, which are not available for the other forecasts.) However, the differences among forecasts, especially that between the Green Book and the MPS, are not large. Sims tests to see whether these differences in performance are statistically significant. He finds strong evidence of Green Book superiority over the SPF model in forecasting inflation at both the one- and the four-quarter horizon, but only weak evidence of its superiority over the MPS model.

Sims also compares the forecasts' performance at predicting growth in output, still using RMSE as the criterion. Here the naïve forecasts do poorly, despite their informational advantage, and the SPF forecasts rank better than they did in forecasting inflation, slightly outperforming the MPS forecasts. Although the Green Book forecasts are still the best at all horizons, their advantage over the MPS and SPF forecasts is quite small and not statistically significant. The Green Book's biggest advantage is at a horizon of three quarters, where its RMSE is 3.2 percentage points, compared with 3.37 and 3.43 percentage points for the SPF and the MPS, respectively.

A second basis for comparison is provided by the correlations of the different forecasts with actual outcomes. The square of the correlation coefficient measures the fraction of the actual variance that is explained by a regression of actual values on forecasts. Whereas the RMSEs are elevated by bias and the failure of forecasts to move proportionately with actual values, a regression with a freely estimated intercept and coefficient removes these two sources of error. If a policymaker were limited to using a single forecast, and used the regression equation to account for bias and scaling error, the forecast with the highest correlation would be the most informative. In the case of inflation, all the forecasts are quite highly correlated with actual values: they explain over 80 percent of the variance in inflation at a horizon of one quarter and over 65 percent of the variance at four quarters. Again the Green Book forecasts perform best, explaining

nearly 80 percent of the variance at four quarters. Sims notes that the Green Book inflation forecast achieves a lower RMSE even though both the MPS and the naïve forecasts have lower bias, and that the Green Book looks even better in the correlation comparisons. As measured by correlations, none of the candidates do as well at forecasting GDP growth as they do at forecasting inflation. But whereas the forecasts did not differ greatly when compared by RMSE, the differences in correlations are more pronounced. The Green Book, which had the lowest RMSE, also has the highest correlation, explaining 17 percent of actual growth four quarters out. The SPF, which had the second-best RMSE, explains only 6 percent at that horizon, and the MPS only 10 percent.

A third way to characterize the value of the forecasts is by their individual contribution in a multivariate regression that includes all forecasts. The Romers, following Ray Fair and Robert Shiller, refer to this sort of regression as measuring the "information content" of forecasts. In the case of inflation, such a regression for four-quarter-ahead forecasts results in a large and significant coefficient on the Green Book forecast and insignificant coefficients on the others (the estimated coefficient on the SPF forecast is negative). For four-quarter-ahead output growth, the results are similar, with a coefficient of 1.4 on the Green Book forecasts and a negative but insignificant coefficient on the MPS forecast. Sims suggests that such equations are useful if interpreted carefully, but he questions whether they should be characterized as providing information content. He illustrates the general point with a model in which the variable of interest is a linear function of one or more forecastable components plus error, and the forecasts are different, noisy estimates of these components. With one forecastable component, the regression coefficients will all be positive, but their magnitudes will depend inversely on the error in the relationship of each to the forecastable component. Sims points out that the difference in relative sizes of the coefficients can be extreme even though the forecasts have very similar forecast error variances. With two forecastable components the interpretation becomes even more problematic. In this case, even though the common component of error is small and very similar among forecasts, the coefficients can be of opposite sign and grow larger as the forecasts become more similar. The coefficients give the optimal weights for a best forecast using a linear combination of the individual forecasts. Those weights have no reliable relationship to any individual forecast's quality.

What accounts for the apparent superiority of the Green Book inflation forecasts? Sims explores three possibilities: better use of time-series data available to all, access to information not publicly available about the Federal Reserve's own likely policy actions, or access to better detailed information about current price developments. These hypotheses correspond to testable restrictions on a vector autoregression (VAR). When the Green Book inflation forecasts are included in a standard quarterly VAR with four lags each of output, the price deflator, the federal funds rate, the commodity price index, and the interest rate on three-year Treasury bonds, they make a substantial contribution to the fit. The coefficients on the Green Book forecast are highly significant in the equations for the deflator and the federal funds rate, even though the VAR includes data not available at the time those forecasts were made. This suggests that the forecast's success is not due to the Federal Reserve making better use of publicly available information.

In the VAR for the federal funds rate, the coefficient on the Green Book forecast of inflation is strongly positive and not significantly different from 1. Sims notes that this is consistent with the view that the Federal Reserve responds to its own forecast of inflation and that its forecasts contribute to the prediction of other variables through their contribution to forecasting interest rates. Further supporting this view is the fact that when the actual value of the federal funds rate is included in the VAR, the evidence that the Green Book inflation forecasts influence the other variables is weak. Sims finds this pattern of results consistent with the idea that the Federal Reserve's advantage in forecasting arises from getting useful information sooner, perhaps including, most importantly, information about its own policy actions.

As already noted, the subjective judgments of experts play an important role in the preparation of forecasts at all of the central banks. The economists Sims interviewed believe the main value of these inputs is in providing a more accurate picture of the current state of the economy, and therefore a more accurate starting point for forecasts, than a model based on a limited set of data available only at monthly or quarterly intervals. They argued that experts pay attention to a large amount of data from disparate sources, some of it nonquantitative, and that they understand the implications of such information for aggregates or for related data that are used in econometric models but that become available only with a

lag. This view of the role of experts fits nicely with the findings discussed above. Sims does not see this role diminishing with time. As the economy changes and new economic problems arise, new data will be needed that cannot be easily or quickly incorporated into an econometric model. He suggests that unusual events are a second reason for using subjective inputs from experts. An oil crisis, an attack on a currency, a September 11, or simply a data collection error—such events are not best treated as simple draws from a historical distribution of random disturbances. Precisely because they are unique or unusual, analysis of such events inevitably involves an element of subjective judgment.

Much of Sims' paper consists of a critical review of the history, current state, and use of statistical modeling for macroeconomic policy. He discusses separately the evolution of central bank models over the last thirty years, the difficulties of incorporating modern developments in theory and econometric practice into policy modeling, and the need to better integrate stochastic modeling with decisionmaking under uncertainty. In addition, he provides a detailed assessment of the primary econometric models currently in use and suggests directions for improvements. This critique is full of insightful and provocative observations that will be only briefly summarized here.

Sims notes that, because econometrics and macroeconomics were active research areas throughout the period, one might have hoped to observe clear progress in moving from the early simultaneous-equations models to the models used today. He concludes, however, that the evolution of modeling technique over this period does not represent progress. He traces this failure to the fact that academic research has largely ignored the difficulties of modeling for macroeconomic policy conducted in real time, and that central bank researchers have tried to adopt modern developments such as rational expectations theory, calibration, and unit-root econometrics, which apply very awkwardly to policy modeling problems. He finds that VAR modeling, with which he himself is closely associated, is also ill suited to policy modeling in important respects.

Sims discusses several features of policy modeling that are awkward both for traditional simultaneous-equations econometrics and for these newer approaches. He concurs with many other practitioners that real-time forecasting and policy analysis require processing a great many variables. This explains why the number of variables in central bank models is large relative to the number of observations and why, as a consequence, statis-

tical methods for dealing with simultaneity do not work well. VARs of the relevant scale have a similar problem. Maintaining a large-scale model requires a large team of researchers and a way of drawing on specialized sectoral knowledge. This leads banks to decentralize modeling and forecasting responsibilities and to allocate responsibility for sectors to individual economists. Although this setup provides useful real-time knowledge, it compromises other properties of the model as a whole. Sectoral experts are often responsible for individual equations of the overall model, and little attention is given to the properties of the resulting system of equations. The result can be a system with poor long-term properties. Sims cites as an example the Federal Reserve's Global model, which ties together the FRB/US model with models for thirty other countries. There is no well-defined solution to the resulting full system, which, besides raising questions of the model's logical consistency, means the model cannot be solved in its "forward-looking" mode.

Sims observes that the theoretically ideal policy model would be derived from a theory of the behavior of economic agents who interact to generate equilibrium. But models that meet this ideal and are practically soluble need to be relatively small and contain only a small number of sources of stochastic disturbance. Such small models do not fit the data nearly as well as less restricted models and are at a severe disadvantage in forecasting. Although this may not concern academic researchers, it has been a major concern of central banks. Their own research efforts have made little progress, however. While acknowledging the inherent difficulties, Sims is quite critical of the compromises made in various central bank models to retain features of forward-looking equilibrium models while attempting to maintain reasonable real-time performance. As he sees it, in trying to incorporate a more rigorous foundation, central banks have given up any serious effort to fit the data. Yet even the improvement in rigor is largely illusory.

For decisionmaking under uncertainty, Sims argues that simple axiomatics as well as the intuition of most decisionmakers leads to thinking of everything that is unknown, including parameter values and even which model is correct, as subject to probability calculations. He is disturbed to find that central bank economists, most of whom have conventional, non-Bayesian training, regard this kind of odds calculation as unscientific. As an example, he cites the question of whether productivity growth accelerated in the 1990s. Researchers he interviewed thought that

if they had proceeded "scientifically," they would have treated the null hypothesis of no change as true until it was rejected at a 5 percent significance level. But the policymakers they worked for wanted to weigh the probability of a change in growth against the cost of erring in either direction—either finding no change when there had in fact been one, or finding a change when in fact there had been none. Sims regards the fact that these elementary applications of Bayesian decision theory were seen as unscientific as an indictment of the way econometrics is now being taught.

Sims provides a detailed discussion of the advantages of a Bayesian perspective on inference. He sees the most persuasive argument for the Bayesian perspective as its ability to make apparently intractable inference problems tractable. That ability, he argues, is increasing. He cites recent work by Frank Smets and Raf Wouter using advanced simulation techniques for Bayesian analysis as the first example of a dynamic stochastic general-equilibrium model that produces a fit to data that is competitive with a Bayesian reduced-form VAR. And he suggests that, with their substantial resources, central banks' research staffs should be able to successfully estimate rigorous dynamic versions of today's central bank primary models. At a more applied level, Sims emphasizes the general applicability of the Bayesian perspective for policymaking. It provides a natural way of taking into account that a model's parameters may drift over time, and it offers a much less complex way for making inferences about unit roots and cointegration than non-Bayesian methods provide. Sims notes that decisionmaking depends on judgment as well as on inference from the data, and the Bayesian perspective formalizes the interaction of the two. When a policymaker confronts results from models that give conflicting forecasts, he or she does not want to be told that no probability weights can be given to the models. But a non-Bayesian has no way of assigning probabilities across several models, all of which fit historical data reasonably well but have different implications for policy. The Bayesian perspective is thus, in Sims' view, the essence of decisionmaking, which weighs uncertain prospects to compare the expected consequences of different courses of action.

The Federal Reserve Act enjoins the U.S. central bank to pursue maximum employment, stable prices, and moderate long-term interest rates. Although low inflation and low unemployment are both thus objectives of policy, economic events often reveal a conflict between them, at least

in the short run. Analytically, this conflict is most commonly embedded in NAIRU or Phillips equations, which relate the current rate of inflation to the gap between the level of output or employment and its full-employment value, and to past or expected rates of inflation. Even economists skeptical of the NAIRU or Phillips relationships are likely to believe that there will be upward pressure on prices when output exceeds its potential and downward pressure when output is significantly below potential. It is therefore natural that rules for the conduct of monetary policy specify that the interest rate on federal funds, the Federal Reserve's principal policy instrument, should depend on the gaps between actual and desired values of inflation and unemployment. Such rules also typically include the "natural rate of interest," the real interest rate that is taken to be consistent with desired rates of employment and inflation. For example, according to the Taylor rule, perhaps the best known of proposed monetary policy rules, the federal funds rate should increase more than one for one with increases in inflation and decrease with a positive employment or output gap. Interestingly, although policy has never been formally guided by any rule, the response coefficients for the Taylor rule that some recommend using to conduct policy provide a reasonably good description of actual behavior of the Federal Reserve for much of the post–World War II period. In the second paper of this issue, Athanasios Orphanides and John Williams argue that uncertainty about key parameters in such policy rules—the natural rate of unemployment and the natural rate of interest—is a major problem. They show that the optimal policy response to unemployment and inflation in conventional rules is significantly altered by such uncertainty, and they propose a particular rule that they find to be robust in the face of it.

The authors first illustrate the problem of uncertainty in the context of a simple Taylor rule that has performed quite well in model-based evaluations. This rule sets the federal funds rate equal to the policymaker's estimate of the natural rate of interest if the unemployment gap (actual unemployment minus the natural rate of unemployment) is zero and inflation is at the policymaker's target. The federal funds rate is then increased by half of any excess of inflation over the target and decreased one for one with the unemployment gap. With this rule a 1-percentage-point overestimate either of the natural rate of interest or of the natural rate of unemployment will result in the funds rate being set 1 percentage point too high, an economically significant error.

The authors consider an elaboration of the Taylor rule with two features that mitigate such errors. First, they include the change in unemployment from the previous period while reducing the weight placed on the unemployment gap and hence the importance of errors in estimates of that gap. Second, they incorporate policy inertia by only partly adjusting the federal funds rate to the perceived natural rate of interest in each period; this dampens the response to errors in those perceptions. This rule nests the simple Taylor rule as one special case and insulates policy from misperceptions of either natural rate as another. In the latter case the rule becomes what the authors call a difference rule, relating the change in the federal funds rate to the inflation gap and the change in unemployment.

How much to insulate policy from errors in estimates of the natural rates depends on the magnitude of those errors. The authors point to abundant evidence that the natural rates vary significantly over time. They cite previous work suggesting that changes in demographics, the efficiency of job matching, productivity, openness to trade, and the incidence of disability and incarceration all influence the natural rate of unemployment, and that fiscal policy and household preferences influence the natural rate of interest. The authors report that other analysts, in estimations that assume constant natural rates, have found confidence intervals of 2 to 4 percentage points for the natural rate of unemployment and 3 to 4 percentage points for the natural rate of interest. Allowing for a stochastically time-varying natural interest rate yields confidence intervals twice that large. And the authors believe that even these underestimate the true uncertainty because they ignore whether the model itself and the way it was estimated are appropriate.

To provide a sense of the importance of these further uncertainties, Orphanides and Williams present estimates over several decades of the time-varying natural rate of unemployment from five previous studies by different authors, along with their own estimates. The latter use both popular univariate filters that simply estimate a time-varying trend, as well as multivariate models in which the unobserved natural rates are assumed to follow a random walk. The authors estimate the natural unemployment rate with two such models. One uses Kalman filter estimation of a NAIRU model relating inflation to the unemployment gap, lagged inflation, and relative oil and nonoil import price inflation. The other is a simple accelerationist Phillips curve estimated using methods proposed by Laurence Ball and Gregory Mankiw. The authors also use a Kalman filter to estimate

a multivariate model for the natural rate of interest, relating the unemployment gap to the real interest rate gap while imposing the condition that inflation is stable if both gaps are zero.

Orphanides and Williams report that, over a wide range of studies, including their own, even "retrospective" estimates—those based on the entire sample period—show significant variation in the natural rates over time. The estimated values of both rates are generally low toward the end of the 1960s, rise through the 1970s, and trend downward thereafter, returning in the late 1990s to levels near those of the late 1960s. There is also significant period-by-period variation across investigators and methods. For the natural rate of unemployment the range of estimates exceeds 2 percentage points in 1960, 1970, and 2000, and 3 percentage points in 1980. Only in the 1990s are the various estimates reasonably close to each other, in a range from 5.8 to 6.4 percent.

From the perspective of policymakers, matters are even worse. Policy has to be based not on retrospective estimates but on "real-time" or "one-sided" estimates, which use only data available when decisions are made. The authors document the differences between real-time and retrospective estimates of both natural rates for their univariate and multivariate models. Not surprisingly, the real-time estimates are noisier, especially in the earlier periods when fewer observations are available. For the early 1980s the variation across models is especially wide, with estimates of the natural rate of unemployment differing by almost 4 percentage points, and those of the natural rate of interest by more than 9 percentage points. The univariate filter estimates are particularly noisy. Furthermore, the real-time estimates of the natural rates contribute little to the accuracy of forecasts of inflation. A one-quarter-ahead inflation forecast using the most successful estimate of the natural rate has a standard error of 1.1 percentage points compared with a standard error of 1.2 percentage points for a simple fourth-order autoregressive forecast. Forecasts of the unemployment rate using estimates of the natural rate of interest are only slightly better than those from a simple second-order autoregression.

The authors use the difference between the real-time and retrospective estimates of the natural rates as a measure of perception error by policymakers. Over the range of models they consider, the standard deviations of these perception errors range from roughly ½ to ¾ of a percentage point for the natural rate of unemployment and from 0.9 to 1.7 percentage points for the natural rate of interest. Within each model, errors for both rates are

highly persistent, typically with serial correlations in excess of 90 percent. The authors again note that these indicators of misperception still significantly understate the true uncertainty confronting policymakers, both because the retrospective estimates are themselves subject to sampling error and because the identity of the true model is uncertain. To give an idea of the importance of model uncertainty, they compute the errors that would be made using any one of their six models' real-time estimates if any one of the other models' retrospective estimates were the true one. Even from this restricted set of models, the range of perception errors increases substantially in this cross-model calculation, with the median standard deviation of the perception error for both natural rates at or above the largest errors from the within-model calculations just described.

Armed with these insights, the authors turn to their key question: how does uncertainty about the natural rates affect economic performance under various monetary policy rules? To answer this question, the authors need a model of inflation and unemployment, including the way both of these respond to the federal funds rate. They also need estimates of the distribution of shocks to the model, including the covariance of those shocks with the perception errors of policymakers using the policy rule. The authors focus most of their attention on a two-equation model that combines forward-looking elements with inertia in both inflation and unemployment. Inflation is a linear function of expected inflation in the next period and last period's inflation, with weights summing to one, and the expected value of the unemployment gap, which is calculated from the retrospective Kalman filter estimate of the natural unemployment rate. In the model, the unemployment gap depends on the expected unemployment gap in the next period, two lags of that gap, and the lagged real interest rate gap calculated from the retrospective Kalman filter estimate of the natural interest rate. For the expected rates of inflation and unemployment the authors use the median forecasts from the Survey of Professional Forecasters. Estimates of this model for the period 1969:1 to 2002:2 give sensible results. The unemployment equation weights expected and lagged inflation about equally. A 1 percent increase in the unemployment gap reduces the rate of inflation by about a third of a percent in the short run. The interest rate gap is statistically significant in the unemployment equation; the response of the unemployment rate to this gap is small in the short run, but in the long run it exceeds 1 percent for a 1 percent change in the interest rate.

Alternative policy rules differ in their vulnerability to policymakers' misperceptions of the natural rates. As discussed earlier, a simple Taylor rule weights such errors heavily, whereas a difference rule is insensitive to such errors. The authors evaluate the performance of these and other rules using several alternative estimates of errors in natural rate perceptions: a benchmark that assumes errors are zero; an error process for the rates implied by the Kalman filter models; and errors that are two and three times the Kalman filter–based estimates. In light of the model and estimation uncertainties discussed above, the authors believe the last of these are probably the most realistic.

The performance of a given policy rule depends on shocks of different sorts, including errors in the equations for inflation and unemployment, innovations in the natural rates, and errors in perception. The covariances as well as the variances of these errors are potentially important. The authors compute the appropriate variance-covariance matrix, finding, among other things, that misperceptions of the real natural interest rate are positively correlated with shocks to the unemployment rate and with misperceptions of the natural rate of unemployment, and that misperceptions of the latter are negatively correlated with shocks to inflation. The authors assume the policymaker has a loss function that is quadratic in the deviation of inflation from its target rate, the deviation of the unemployment rate from its natural rate, and the change in the short-run interest rate. They use this loss function in evaluating the performance of several policy rules including the original Taylor rule, the generalized rule discussed earlier that allows partial adjustment of the federal funds rate and includes the change in unemployment, and their difference rule.

The authors first compare the performance of Taylor rules specified with various parameter values, including those originally suggested by Taylor, those that minimize loss in the absence of perception error, those that minimize loss using the perception errors from a correctly specified Kalman filter model, and those that minimize loss using errors that are two and three times that size. For each of these parameter sets, they then compute the expected loss for each of the alternative assumptions about perception error. When perception errors are zero, the original Taylor rule performs quite well, nearly as well as when the rule parameters are optimized for this environment. The optimal parameters in the absence of perception errors place much greater weight on the unemployment gap than the original Taylor rule and less weight on inflation. However, with

these parameters the Taylor rule is highly vulnerable to errors in perception. It is least vulnerable to such errors when set to its original parameter values, because these are the least responsive to the output gap. Yet even with these parameters, the loss when errors are from the Kalman filter estimates is roughly double what it is when errors are zero. With Kalman filter errors, the loss using the Taylor rule optimized for no errors is almost triple the loss of the original Taylor rule. In contrast, the rule optimized for the Kalman filter errors pays only a small penalty if in fact errors are zero, and it is much less vulnerable to perception errors two or three times as large as those from the Kalman filter estimates.

These results suggest an important asymmetry. In the authors' examples the cost of underestimating misperceptions is significantly greater than the cost of overestimating them. Policy rules optimized on the assumption that perception errors are small are characterized by large responses to the unemployment gap, which can be very costly. Overestimating perception errors, on the other hand, leads to timid responses, but this is not very costly. Several other generalizations are apparent from the authors' results. Although the weight assigned to interest rate variability in the loss function is small, it contributes noticeably to loss because the rule causes interest rates to vary a great deal. The standard deviation of the quarterly change in rates is greater than 2.8 percentage points in the first three cases, and it exceeds 4 percentage points for the rule optimized for the Kalman filter perception errors. As the degree of uncertainty rises, the optimal weight on the unemployment gap shrinks toward zero, and the weight on inflation gradually increases. Finally, at any of the parameter values considered, the difference rule dominates simple Taylor rules unless perception errors are in fact very small. This finding leads the authors to favor this rule as "robust."

The authors go on to consider the performance of a more flexible, generalized rule that allows various amounts of policy inertia and introduces changes in unemployment as an additional variable. In seeking the optimal parameters for this rule, they find that, as the degree of misperception rises, the optimal amount of inertia and the weight on the change in unemployment increase rapidly whereas the weight on the unemployment gap rapidly declines. The weight on inflation varies relatively little. When the generalized rule is optimized for the uncertainty associated with the Kalman filter model, it performs quite well over the full range of perception errors. It does modestly better than the robust rule for errors less than

one and a half times the Kalman filter errors, and modestly worse for larger errors.

The authors recognize that their results may be sensitive to their specification of the economic model determining inflation and unemployment. To explore this sensitivity, they perform their experiments under two alternative specifications. One is a "new synthesis" model that is entirely forward looking, with no lagged values for inflation or unemployment. The other is a backward-looking "accelerationist" model with no expectational variables. They find that the robust rule does quite well in the new synthesis model, generally outperforming the generalized Taylor rule. However, the robust rule does not do well in the backward-looking accelerationist world: it performs substantially worse than the generalized Taylor rule except at the highest levels of misperception. Given the economic profession's lack of agreement on how to specify the economic models used for policy analysis, it is disconcerting that the authors' conclusions are so sensitive to the choice of model.

The authors' experiments indicate that misperceptions of the natural rates can be quite costly, particularly if policymakers are overconfident about their knowledge of those rates. The costs of overconfidence are likely to be greatest when the natural rates change significantly, so that the errors in policymakers' real-time estimates are likely to be large. The stagflationary period of the 1970s and the boom of the 1990s are two periods when, in retrospect, large shifts in the natural rate of unemployment apparently occurred. The authors document that contemporary observers were slow to adjust their beliefs in both those periods. They cite Herbert Stein, Paul McCracken, and Arthur Burns as distinguished economic policymakers who, looking back, recognized their overoptimism in the 1970s. And they show that, in the 1990s, most professional forecasters, as well as the Congressional Budget Office and the Council of Economic Advisers, only belatedly reduced their estimates of the natural rate in the 1990s.

To see what difference the design of policy can make in episodes like these, the authors compare the performance of the Taylor, generalized, and robust rules for stylized versions of the two periods. They assume that in the 1970s both natural rates increased by 1.5 percentage points over a two-and-a-half-year period, at the beginning of which policymakers knew the true levels of both rates. Given the speed of adjustment implied by the authors' Kalman filter model, errors in estimates of the natural rates grow

to approximately 1 percentage point and dissipate slowly over many years. Following the original Taylor rule in the presence of these perception errors results in persistent inflation, as the policymaker strives for an extended period to achieve what is in fact too low a level of unemployment. The revised Taylor rule, which increases the policy response to the unemployment gap, does worse. The generalized rule, optimized for the optimistic view that there are no perception errors, does little better. However, the robust rule, which ignores estimates of the unemployment gap, results in much less persistent inflation with a very minor cost in increased unemployment. For the purposes of the authors' experiment, the 1990s are treated as the mirror image of the 1970s, and the results are a mirror image of the 1970s results as well. Following the generalized rule, again optimized on the assumption of no errors, results in deflation, with inflation falling by almost 6 percentage points and staying well below its initial value for many years. By contrast, the robust rule produces an outcome much like the "Goldilocks" experience that actually occurred. The authors thus see the contrast between actual policy behavior in these two periods as illustrating and supporting their analytical results favoring their difference rule.

THE NATIONS OF EUROPE HAVE become increasingly integrated during the past decade, first through strengthening trade and investment ties within the European Union and then through the buildup to and eventually the adoption of the euro as a common currency. Many of the promised benefits of this greater integration have been apparent for years. But the economic stabilization problems that had worried many have recently emerged and become a source of controversy. In the third paper of this issue, Olivier Blanchard and Francesco Giavazzi call attention to a development that was less widely anticipated and that has thus far received little notice, namely, the growing dispersion in the size of current account balances among European nations.

As a formal backdrop for their empirical work, Blanchard and Giavazzi present a traditional two-period model of how current account positions evolve, which they adapt to the particular circumstances of emerging monetary union among countries with very different initial positions. In the model, saving and investment in each country are determined by households allocating consumption over time subject to an intertemporal budget constraint. In each period, consumption in all coun-

tries together equals production in all countries, and for each country the budget constraint and a utility function determine consumption over the two periods. For this setup the authors show that the current account balance for each country depends negatively on the growth of its real income relative to growth in world income, negatively on the interest rate the country confronts relative to interest rates paid by other countries, and positively on the elasticity of the demand for the country's goods. This simple model captures what the authors see as the principal mechanisms through which greater economic integration, particularly monetary union, would affect current accounts. They assume that a less advanced economy entering the union will grow faster than the average economy, as its convergence toward the income level of the more advanced economies is hastened by joining the union; that it would find its interest rate falling toward rates in other economies as exchange rate risk is eliminated; and that it would find that greater trade integration increases the price elasticity of demand for the goods it produces, making it easier to repay current borrowing with future trade surpluses. The model predicts that each of these main developments pushes the less advanced economy's current account into deficit.

The authors assess recent developments in Europe against these predictions. They first examine panel data for three successively more focused country groupings. The first, which they call OECD minus, includes twenty-two of the thirty countries of the Organization for Economic Cooperation and Development, including all the European members except Luxembourg, Turkey, and the Central European countries. (Luxembourg is an unusual case, with reported current account surpluses consistently around 30 percent of GDP; Turkey is much less developed than the other European OECD countries; and the Central European countries lack adequate data.) The second grouping, consisting of fourteen countries, is the European Union, again excluding Luxembourg. The third is the euro area, consisting of the twelve countries that have adopted the euro, again excluding Luxembourg. A fourth grouping, which the authors calls euro minus, omits Greece and Portugal to test whether results for the euro area arise entirely from the inclusion of these two poorest members. The current account data are from the European Commission's Annual Macroeconomic Database (AMECO) and are based on countries' national income accounts. The authors report that their results are much the same when the alternative data set published by the OECD is used.

The authors' model suggests that increased integration results in greater cross-country variation in current accounts. They find that annual cross-country standard deviations of current accounts for 1975–2000 rise steadily for all four groupings after the mid-1980s. Before that, in the early 1980s, a sharp rise in the standard deviation is seen, followed by a decline. This earlier rise and fall is driven by large deficits in Ireland and Portugal, which Blanchard and Giavazzi argue arose from temporary shocks and were unsustainable. They therefore first focus their attention on the years starting in 1985. For two subperiods, 1985–93 and 1994–2000, they estimate cross-country regressions of current account balances on income per capita. For each country grouping the coefficient is positive in each period, and larger in the second, consistent with their model's prediction that with integration poorer countries should run current account deficits. The prediction is also supported by panel regressions for the entire 1975–2001 period that, for each year and country, relate the current account–to–GDP ratio to income per capita relative to that in the entire group, and to a common time effect and two variables that control for effects of the dependency ratio and cyclical movements in output in each country. The effect of relative income per capita is allowed to vary over time, providing a test of whether integration strengthens the effect of relative income on current accounts. With this specification, the authors find a strengthening effect after the mid-1980s for all groupings except OECD minus. The relation appears equally strong for the euro area whether or not Portugal and Greece are included. By the mid-1990s the coefficient in regressions using data for the EU group is statistically and economically significant. At its 2000 value of 0.2, this coefficient indicates that a country with an income per capita 40 percent below the EU average, roughly the position of Portugal and Greece, should have a current account ratio 8 percentage points below the EU average. For the euro area the effect is even larger.

Some alternative specifications support these basic findings. Unless Ricardian equivalence holds, the model predicts that public saving, like household saving, should be positively related to current account balances. Adding to the regressions the ratio of structural primary government budget balances to GDP, the authors find a significant positive relation for all groupings. For the OECD grouping, a 1-percentage-point rise in the budget balance ratio leads to a 0.2-percentage-point increase in the current account ratio. The previously estimated effect of income per capita is lit-

tle changed by adding government budget balances to the regressions, indicating that the increased divergence in current accounts is not explained by a growing divergence in public saving across countries. Similarly, adding the Central European countries to the sample (useful data for which begin only in the early 1990s) also confirms the main results. Finally, the authors add M3, a broad measure of money, to their current account regressions as an additional control variable that proxies for the borrowing opportunities available to domestic residents and firms. They find that M3 is strongly negatively related to current account balances, and they infer that internal financial liberalization plays an independent role in determining the distribution of current account balances across countries.

The authors next address whether saving or investment is the main channel through which the current account is affected. To do this they alternately substitute first the saving-to-GDP ratio and then the investment-to-GDP ratio for the current account–to–GDP ratio as the dependent variable in their basic pooled regression. In regressions explaining the saving-to-GDP ratio, the coefficient on relative income per capita is, after the mid-1980s, positive with a rising trend for the euro area, positive but with a weaker trend for the European Union, and neither significant nor showing a trend for the OECD. In regressions explaining the investment-to-GDP ratio, the sign on income is generally negative, as the model predicts, but the authors find no well-defined trend in the coefficient. In all groupings some downward trend begins by 1995, but there is severe volatility before that. The authors conclude that the increased dependence of the current account on income per capita reflects mainly effects operating through saving rather than through investment.

To explore the mechanisms behind these panel data results, Blanchard and Giavazzi look more closely at the experience of Portugal and Greece. To minimize the effects of cyclical factors on the current account, they choose 1985–91 as the base period for Portugal and compare performance then with performance in 1996–2001, a period when integration was well advanced. They first calculate changes between these periods in the average values of the current account, investment, and several components of saving, all expressed as ratios to GDP. The current account balance declined by 10.6 percentage points of GDP, reflected in a 2.8-percentage-point increase in investment and a 7.8-percentage-point decline in saving. In allocating this saving decline, the authors adjust the government surplus

for inflation by adding an amount equal to the inflation rate times the government debt denominated in domestic currency (roughly half of total government debt in all years). On the assumption that this portion of the debt is entirely owned by domestic households, they subtract the same amount from household saving. With these adjustments, which move over 5 percentage points of the reduction in saving from households to the government, 2.2 percentage points of the decline in saving came from a reduction in the government's surplus, 3.8 percentage points from a decline in household saving, and 1.8 percentage points from a decline in corporate saving.

The authors see increased borrowing by Portuguese households as a key factor in the decline in household saving. From 1995 to the end of 2001, household debt, mainly in the form of mortgages and consumer loans, increased from 40 percent to 93 percent of GDP. Finding no evidence of substantive changes in the types of financial products available over this period, the authors conclude that falling interest rates are the main factor behind this rise in indebtedness. Between 1992 and the end of the period, annual short-term interest rates fell from 16 percent to 4 percent, and real short-term rates fell from 6 percent to near zero. The authors attribute much of this decline to financial integration, which eliminated country risk and opened the euro interbank loan market to subsidiaries of Portuguese banks. In 2000 the net increase in foreign indebtedness of Portuguese banks amounted to 10.7 percent of GDP, which is more than the current account deficit in that year and much larger than all other portfolio and investment flows.

Net foreign direct investment, which had been an important source of capital inflows in the mid-1980s when Portugal entered the European Union, turned negative by the mid-1990s. Since then inflows have increased rapidly, but outflows have increased even more. The authors provide an interesting interpretation of these developments. Outflows have mainly taken the form of acquisitions, with over 40 percent going to Brazil. Portugal's bankers presumably have a comparative advantage in understanding and dealing with Brazilian firms. When this long-standing special relation was coupled with Portugal's newfound ability to raise funds in the euro area capital market, it led to substantial European direct investment in Brazil being carried out through Portugal, giving rise to both inflows and outflows of capital.

For their analysis of Greece, Blanchard and Giavazzi use 1981–91 as their base period and compare it with 1996–2001, omitting the intervening

years to avoid cyclical distortions. Between the two periods, the ratio of the current account to GDP fell by 3.5 percentage points; all of this decline is attributed to a decrease in saving, which was also the most important factor in Portugal's experience. However, the allocation of the saving decline in Greece is very different from that in Portugal. Adjusting the government deficit for inflation, which again makes an important difference, and allocating the adjustment to households as before, the authors find that the private saving ratio fell by 7.7 percentage points while the government saving ratio rose by 4.3 percentage points. The available data do not permit an allocation of private saving between households and corporations before 1995. Between then and 2000 there was little change in the household saving ratio, while the corporate saving ratio fell by 4.3 percentage points.

Getting behind these numbers, the authors report that the decline in corporate saving reflected a shift in the financing of firms from internal finance to share issues: capital raised in the stock market rose from zero in the mid-1990s to 8 percent of GDP in 2001. Over the same period the outstanding volume of consumer loans as a fraction of GDP rose from 1.6 percent to 6 percent, and the volume of mortgage loans rose from 4.5 percent to 12 percent. The authors note that, in contrast to what happened in Portugal, domestic financial liberalization facilitated this increase in borrowing: before 1997, consumer loans were virtually prohibited in Greece. Outward foreign direct investment has been about 3 percent of GDP, with about one-third going to the other Balkan and Mediterranean countries. Direct investment in Greece from abroad has been slightly larger, resulting in small but positive net foreign direct investment into Greece in recent years. The authors find evidence that the move to the euro has been important in making financial investments in Greece more attractive: purchase of government bonds by foreigners has become one of the main capital flows financing the current account deficit. The authors also report that unit labor costs in both Greece and Portugal were little changed between 1995 and 2001, and they conclude that changes in competitiveness were not an important factor affecting the current account in either country.

Although the authors thus see developments in Greece and Portugal as broadly supporting both their stylized model and their findings based on panel data, they recognize that Ireland, which also used to be much poorer than the other EU countries, has had a different experience. They note

that Ireland, whose experience Patrick Honohan and Brendan Walsh analyzed in the Spring 2002 issue of the *Brookings Papers,* has enjoyed very rapid GDP growth, which generated a large increase in government saving. And they see this swing in the fiscal balance as the main factor accounting for the current account surpluses that Ireland has experienced.

In a provocative 1980 paper, Martin Feldstein and Charles Horioka found that investment and saving within a country are highly correlated over time. Such a correlation would be expected in a relatively closed economy but might be expected to decrease with the opening of capital markets. Blanchard and Giavazzi revisit this issue by regressing investment-to-GDP ratios on saving-to-GDP ratios for their four country groupings and for the OECD as a whole (again omitting Luxembourg). The regressions are estimated for the period 1975–2000 and allow for time variation both in the intercept and in the key coefficient relating saving and investment in each country. For all their country groupings, they find that the coefficient relating saving to investment varies noticeably over time and clearly declines in the later years of the sample, indicating that investment and saving are becoming less correlated. For the European Union and the euro area, the coefficient has declined to zero or below by the end of the sample.

The authors conclude by discussing whether the current account developments that they have analyzed should be a source of concern for Europe's policymakers. The most relevant of these concerns is how countries will have to adjust in the future to service or repay the capital flows that correspond to present current account deficits. The need for eventual trade surpluses will in due course require either a real depreciation or a decline in domestic aggregate demand. With no scope within the euro area for countries to pursue an independent monetary policy or to revalue their currencies in nominal terms, this implies some combination of tighter fiscal policy and inflation to worsen the terms of trade. The authors reason that theory offers little guidance about when fiscal policy should be adjusted. However, they argue that there is no case for tightening fiscal policy now to eliminate the present current account deficits. Their model suggests that these deficits reflect a move toward optimal allocation of consumption and investment over time. For this reason they view benign neglect of today's current account deficits as a reasonable course of action.

After two decades during which it had been disappointingly slow, productivity growth in the United States quickened dramatically during the

second half of the 1990s. Because productivity growth varies substantially over both short and medium horizons, economists were slow to recognize this acceleration. But by now it is widely accepted, and attention has turned to understanding its causes. In the final paper of this issue, William Nordhaus sheds some new light on this question using data on value added by industry to construct measures of labor productivity that differ from those commonly used by analysts. With this data set, which provides measures of output derived from the income rather than the product side of the national accounts, he estimates productivity growth for the economy as a whole, for the nonfarm business sector, and for that part of the economy where real output is well measured. He also calculates the effects of changing industry composition on measured productivity, and he creates a new aggregate consisting of those sectors that include the "new economy," to allow an assessment of the contribution of the new economy to overall productivity.

The income-side data that Nordhaus uses are available by industry from 1977 through 2000. Most of his analysis breaks the period into three parts: 1977–89, 1989–95, and 1995–2000, the last being the period when productivity accelerated sharply according to all indicators. Although Nordhaus makes no attempt to adjust the data for the business cycle, the breakpoints he chooses avoid any extremes of the cycle. Using data on nominal output, real output, prices, and hours of work from the Bureau of Economic Analysis, he assembles the industry data into his new economy and well-measured output aggregates and forms an income-side measure for the nonfarm business sector to compare with the widely used Bureau of Labor Statistics (BLS) measure of nonfarm business productivity based on product-side data. The largest difference between these two measures occurs in the 1995–2000 period, when the income-side estimates grew noticeably faster than the product-side estimates. As a result of this large discrepancy in output estimates, and a smaller difference in estimates of hours worked, productivity grew 0.4 percentage point a year faster in Nordhaus's estimates than in the official BLS data. For the earlier years the two estimates of productivity growth are much closer.

Nordhaus's well-measured output omits those sectors of the economy for which he judges output is measured too poorly to be useful in analyzing productivity growth. In some of these sectors, such as government and education, the official data make no serious attempt to measure output and instead proxy it by an index of inputs such as employment. In other

sectors the official output data are based on seriously defective techniques for deflating nominal magnitudes. In construction, insurance, and banking, the price indexes used for deflation are not representative of the range of outputs in the sector; in other sectors, which Nordhaus has identified in discussions with experts inside and outside of the Bureau of Economic Analysis, the price indexes do not adequately capture quality change or the introduction of new goods and services.

Taking all these issues into account, Nordhaus excludes the following major sectors from well-measured output: government and government enterprises; construction; finance, insurance, and real estate; and services other than software, other business services, hotels, and repair services. The remaining, well-measured sectors accounted for 57 percent of nominal GDP in 1977 and 50 percent in 2000. By comparison, nonfarm business output, the aggregate most commonly used in analyzing productivity, remained about 75 percent of nominal GDP throughout 1977–2000. Annual productivity growth for well-measured output averaged 2.0 percent over 1977–89, noticeably faster than the 1.3 percent average for (income-side) nonfarm business. Between 1977–89 and 1995–2000 there was a 1.3-percentage-point speedup of productivity growth in well-measured output, somewhat smaller than the 1.6-percentage-point speedup for the nonfarm business sector.

To analyze these developments further, Nordhaus starts from the observation that aggregate productivity growth reflects both changes in productivity within sectors and changes in the composition of sectors within aggregate output. To identify these sources of change, he breaks down productivity growth in his income-side aggregates into three components: a "pure productivity effect," a "Baumol effect," and a "Denison effect." Now that the National Income and Product Accounts use chain-weighted, or Fisher, indexes to measure output and prices over time, these components sum approximately to overall growth in labor productivity, omitting only small interaction and second-order effects. The pure productivity effect is given by the sum of productivity growth rates of different industries weighted by their nominal output shares in a base year. It can be thought of as the change in productivity that would occur if each industry's share of nominal output did not change. The Baumol effect, named for William Baumol, who pioneered work on "unbalanced growth," identifies the effect of changing industry output shares on overall productivity growth; it is given by the sum of industry growth rates weighted by the departure over

time of nominal output shares from shares in the base year. If industries with slow productivity growth account for a rising share of nominal output over time, as Baumol first reported, this in itself would reduce aggregate productivity growth below that measured by the pure productivity effect. Finally, the Denison effect is given by the sum of growth rates in industry labor inputs weighted by the difference between the industries' shares of nominal income and labor input. It captures the change in productivity that arises from changing shares across industries with different levels of productivity. This effect is named for Edward Denison, who first demonstrated that the movement of workers from low-productivity agriculture to high-productivity industry accounted for an important part of postwar growth in Europe and elsewhere.

For well-measured output, the speedup in the pure productivity effect after 1995 is slightly larger than the total speedup in productivity growth, indicating that none of the acceleration in productivity resulted, on balance, from sectoral reallocations. Also, annual growth from the pure productivity effect averaged about 0.2 percentage point more than total productivity growth throughout the 1977–2000 period. The results were much the same for the entire nonfarm business sector. Interestingly, in both the well-measured sectors and the nonfarm business sector, Denison effects (plus small residuals) subtracted from productivity growth in all three periods. Baumol effects alone added somewhat to productivity growth in 1977–89 and were negligible thereafter. In Nordhaus's words, "Baumol's cost disease has been cured, or at least is in remission."

Nordhaus identifies four sectors where elements of the new economy are concentrated: electronic and other electric equipment (which includes communications equipment and semiconductors), industrial machinery and equipment (which includes computer hardware), telephone and telegraph (telecommunications services), and software. In 1995–2000 the first two sectors recorded impressive accelerations in their own productivity, while the other two had healthy but unspectacular productivity growth. To measure the contribution of each sector to the acceleration of aggregate productivity, Nordhaus constructs Fisher indexes of output and labor input both for aggregates that include these industries and for properly scaled aggregates that omit them. The difference in productivity growth between them provides his measure of the contribution of the new economy to productivity growth. Nordhaus's analysis is thus limited to the direct contribution of these industries through their own productivity growth and their

changing contribution to the total production of goods and services. For reasons discussed below, he does not attempt to measure the contribution of new economy goods and services to productivity growth in other industries through their application of new economy capital.

Nordhaus then applies his method of isolating the contribution of new economy sectors to his three economic aggregates: the total economy (gross national income), the nonfarm business sector, and well-measured output. In each of these he finds that the new economy sectors raised productivity growth noticeably in each subperiod. However, he finds they accounted for less of the productivity acceleration between 1977–95 and 1995–2000 than some other analysts have found. In well-measured output, productivity growth quickened by 1.3 percentage points; the new economy sectors accounted for 0.6 percentage point of this acceleration. In the nonfarm business sector, productivity growth accelerated by 1.6 percentage points, with only 0.3 percentage point accounted for by the new economy. And in the economy as a whole, productivity quickened by 1.0 percentage point, with the new economy accounting for 0.3 percentage point. The difference between total productivity growth and the part of it attributed to the new economy shows that a substantial acceleration of productivity occurred in industries outside the new economy. Among individual industries, Nordhaus finds that three of the new economy industries were among the top ten in their direct contribution to the productivity acceleration. However, the three largest contributors were retail trade, securities and commodities brokers, and wholesale trade—none of which is part of the new economy as conventionally understood. Together these three accounted for 1.05 percentage points of the 1.61-percentage-point acceleration in the annual rate of productivity growth in nonfarm business.

Nordhaus recognizes that the total contribution of the new economy to labor productivity includes the enhancement of productivity through the application of new economy capital in industries that use that capital. And he cites a number of studies that have recently found a very large role for capital deepening in explaining the acceleration of productivity. In their comments on the paper, both Robert Gordon and Daniel Sichel provide informative reconciliations of Nordhaus's new findings with their own updated results that incorporate estimates of the contribution of capital stocks to productivity growth using conventional product-side data. However, Nordhaus offers some important reservations about such calcu-

lations, particularly when applied to the 1995–2000 period. In those years, short-lived high-technology capital, for which there is little historical evidence to gauge service flows, dominated the data on the growth of capital services; the user cost of capital calculated from interest rates in these models was exaggerated because it ignored the cheap finance available from high equity prices; and the stock market boom led to unproductive overinvestment in some industries. Using pooled regressions that explain sectoral output growth with growth of both net capital stocks and labor, he buttresses these misgivings by showing that the role of capital stocks in explaining sectoral output is not well identified.