Editors' Summary

THE BROOKINGS PANEL ON Economic Activity held its seventy-fifth conference in Washington, D.C., on March 27 and 28, 2003. This issue of *Brookings Papers on Economic Activity* includes the papers and discussions presented at the conference. The first paper explores why some emerging market economies are prone to fall into financial crisis at levels of external indebtedness that more advanced economies, and even some other developing economies, seem able to manage. The second paper reviews the current U.S. fiscal situation against the historical record and finds the present is so different that the past is an unreliable guide to how either the economy or future policy will respond. The third paper offers a theoretical analysis of optimal monetary policy in the face of a liquidity trap, with a focus on the importance of expectations. The fourth paper discusses a new methodological approach to economic policymaking under uncertainty, with particular attention to uncertainty about economic models. The concluding report analyzes whether new rules for corporate pension accounting promulgated in the 1980s misled investors into overvaluing the stocks of firms with pension plans during the 1990s market boom.

DEVELOPING ECONOMIES HAVE BEEN vulnerable to financial crises for a long time and under a wide range of exchange rate regimes and international financial structures. Although economists have used theoretical models to explore the link between a country's external debt and its vulnerability to crisis, little empirical work has been done to quantify that link or to identify what makes some countries more vulnerable than others. In the first paper of this volume, Carmen Reinhart, Kenneth Rogoff, and Miguel Savastano draw on the experience of a large number of countries at different stages of development to address these issues. Their perspective is informed by a number of observations about individual

countries' ability to borrow. Advanced industrial countries typically have access to international capital markets irrespective of their debt burdens. The poorest countries, by contrast, are simply shut out of these markets. The large group of middle-income countries find that their access varies with their economic situation. The authors introduce the concept of "debt intolerance" to describe the fact that some countries are more likely than others to get into financial trouble from taking on a given amount of debt. They also attempt to quantify debt intolerance and the factors that lead to it.

Sovereign defaults have a long history. The authors record sixteen defaults by European nations between 1501 and 1800, a period for which data are probably incomplete, and no fewer than thirty European defaults in the nineteenth century. Focusing on present-day emerging market economies, they show a close association between a country's recent *Institutional Investor* rating (IIR), which reflects the views of economists and analysts from the financial community about a country's creditworthiness, and its long-term record with respect to default or debt restructuring. (Calculations of country risk based on market yield spreads give rankings similar to the IIR, but data on yields for many countries' debt are available only for recent years.) Focusing on a sample of emerging market economies, and using data that go back at least a few decades and to 1824 for most, the authors show that those with at least one external default or restructuring had an average IIR of 42 in 2002, whereas those with no history of default had an average rating of 62 that year. By comparison, the ratings of six advanced economies with no history of default average 89. The authors also show a large difference in inflation experience across these groups. Countries in the lowest IIR category in 2002 spent, on average, one quarter of the past forty-two years with annual inflation exceeding 40 percent. The authors note that most of the countries with some history of default are "serial defaulters." Most have spent roughly a quarter of the time for which the authors have data in a state of default or restructuring, and Mexico has been in this condition nearly half the time.

The ability to borrow freely is a hallmark of developed financial markets, which help channel funds to efficient uses. And default in the aftermath of a negative shock can occur even in a well-functioning world financial system. Whether the costs of default exceed the benefits of borrowing is a key question. The authors do not attempt such an evaluation, but they believe that, for many emerging market economies, any potential

benefits from borrowing as freely as they have are outweighed by the direct and collateral costs of default. They note that external borrowing has often been driven by shortsighted governments eager to boost consumption rather than to fund high-return projects. More important, they argue that default imposes large systemic costs through the damage it does to the country's banking and financial system, and that, having once borne those costs, a country is typically more willing to risk default again, further weakening those institutions. Some countries take great pains to avoid this vicious cycle by remaining in the no-default group even when they find they have overborrowed. And the authors reason that those countries that have not resisted default are likely to be exceptionally debt intolerant in that borrowing levels that might be acceptable for others are dangerous to them.

How much debt is too much? Data covering the past thirty years for middle-income countries, which include most emerging market economies with access to private external financing, show a wide range of indebtedness among defaulters. External debt in these countries at the time of their default averaged 71 percent of GNP and ranged from 13 percent (Russia in 1991) to 214 percent (Guyana in 1982), with a median around 60 percent. In 13 percent of the defaults, external debt was below 40 percent of GNP, and in 13 percent it was greater than GNP. Despite this wide range, the authors show that the defaulters, as a group, look substantially different from the nondefaulters. Averaging each country's data over 1970–2001, they find that the frequency distribution of defaulters is massed around a higher ratio of debt to GNP. Since 1979, when IIR scores first became available, individual countries' IIRs have varied widely over time, and there has been a positive correlation between risk (as measured by subtracting the IIR from 100) and debt-to-GNP ratios.

As a way of structuring their ideas, the authors sort countries into borrowers' "clubs" according to their debt intolerance, which the authors measure by a combination of IIRs and debt ratios. They first sort according to whether a country's IIR is within 1 standard deviation of the mean IIR of the sample over 1979–2002, or above or below that interval. Countries with a mean IIR above 67.7 are in club A; these countries are all advanced economies and enjoy virtually continuous access to capital markets, whatever their debt ratios. Countries with a mean IIR below 24.1 are in club C; these countries rely primarily on grants and official loans for their external finance. Countries with mean IIRs between these values are

in the middle group, club B, and have uncertain access to capital markets. This is the group on which the authors focus their analysis. They further sort the countries in club B into four regions according to whether their IIR is above or below the mean for all countries, and whether their debt-to-GNP ratio is above or below 35 percent, the mean ratio for non-defaulters. The authors posit that countries in the lowest region, region IV, are at risk of falling into club C, whereas countries in region I have a chance of graduating into club A.

The authors turn next to what it takes to bring about such movement. Although they recognize that a range of political and structural changes could affect a country's status as a borrower, they focus on what might be inferred from its debt and economic history alone. They perform cross-sectional regressions explaining countries' average IIR during 1979–2000 with a number of variables that measure default history and inflation experience (both of which are largely predetermined in their sample) and average debt ratios. They also use dummy variables to separate countries in club A from the others. Estimates are presented both for ordinary least squares (OLS) regressions and for regressions using the average debt ratio during 1970–78 as an instrument, to account for endogeneity in the debt ratio. The estimated coefficients are uniformly of the sign the authors expect. The debt-to-GNP ratio has a strong negative effect on the IIR for countries in clubs B and C, but the relationship is positive for club A countries (where total rather than external debt is used as the debt measure). The percentage of time since 1948 that a country has spent with annual inflation over 40 percent has a negative effect on the IIR, as does the percentage of years since 1824 spent in a state of default (or since 1946, but the statistical significance of this result is lower). The number of years since the last default has a positive effect, indicating that markets gradually, if only partially, forgive. (These results generally pass conventional significance tests in the OLS regressions but generally do not in the instrumental variables regressions.) The authors also report on panel regressions that relate the IIR to debt ratios and dummies separating the data into three periods. In these regressions (which again span 1979–2000), the pre-debt crisis years (before 1982), the years of the debt crisis and later of debt relief under the Brady Plan (1983–93), and the postcrisis years (1994–2000) appear as distinct periods, with IIRs substantially lower at given debt ratios in the middle period. The authors interpret these results as showing that debt intolerance is long lived, with

IIRs high relative to debt ratios before the early-1980s debt crisis, then falling sharply as the crisis unfolded, and remaining low for a decade thereafter.

Turning to the implications for the behavior of individual countries in club B, the authors use their preferred cross-country regression to ask how a country's IIR would vary with its debt ratio, holding other things constant. To control for the nondebt regressors, they multiply their actual values for each country by their estimated coefficients and calculate predicted values of each country's IIR for varying debt ratios. They illustrate their findings by comparing Argentina and Malaysia. Argentina stays in region I, the least debt intolerant of the regions in club B, only as long as its debt-to-GNP ratio does not exceed 15 percent, and it falls into region III at debt-to-GNP ratios between 15 and 35 percent. Malaysia, with its better history of default and inflation, remains in region I so long as its debt ratio is below 35 percent. Both countries fall into region IV when their debt-to-GNP ratio reaches 40 percent. The authors also use Brazil's experience to show how their estimated equation predicts IIRs, and they use these predicted IIRs together with actual debt ratios to calculate a predicted region. They then compare this predicted region with that assigned using actual IIRs and debt ratios. Brazil was in default or restructuring from 1983 through 1993. Before that, its assigned region was better than its predicted region, indicating that the actual IIRs underestimated the country's riskiness in this period. From 1993 to 2000, both assigned and predicted values are in agreement, placing Brazil in either region III or region IV.

The authors search for candidates for graduating out of club B into club A by looking for those countries whose assigned region within club B has been well above the country's predicted region in recent years. The idea is that, as a country approaches club A status, the coefficients based on the club B regressions would persistently overestimate its debt intolerance. During 1992–2000 the five countries with the largest gaps between their actual and their predicted region were, in descending order, Greece, Portugal, Thailand, Malaysia, and Chile (the last three with substantially smaller gaps than the first two).

If its borrowing were used for highly productive investment, a country might be able to grow out of a high debt ratio. For the period 1970–2000 the authors identify twenty-two episodes in which middle-income countries reduced their debt-to-GNP ratios by at least 25 percentage points

within a three-year period. This is a high hurdle that countries would find difficult to clear through growth alone. The authors find that fifteen of their twenty-two episodes coincided with a default or a restructuring and that in only four of these was growth the principal factor in the decline of the debt ratio. It was an important secondary factor in seven other cases. In the seven episodes that involved no default or restructuring, rapid growth was the main factor in two and an important secondary factor in four others. The authors observe that most of the seventeen countries that took part in Brady bond restructurings in the 1990s were not on the list of episodes they examined, because their debt ratios did not decline enough. Indeed, the debt ratios for many of these countries have risen above their levels at the start of their Brady deals, and three countries that made use of these restructurings have again defaulted. The authors suggest that future restructurings should be designed so as to reduce the likelihood of such relapses.

Historically, external debt has been the main or only source of financing fiscal deficits for emerging market economies, and the authors focus on it in most of their empirical analysis. However, they also note that domestically issued debt is becoming increasingly important for many middle-income countries, and they warn that it now subjects these countries to many of the same perils to which foreign debt has subjected them in the past. In some, domestic debt has already been at the heart of problems. Mexico's *tesobonos* were dollar-denominated domestic securities and therefore exposed the country to exchange rate risk, and similar debt has contributed to recent crises in Argentina and Uruguay. The authors acknowledge that even dollarized domestic debt is not identical to foreign debt in its characteristics and implications: most domestic debt is owed to domestic residents, and there is less risk of a cutoff of international trade credits in the case of a domestic default or restructuring. But they also argue that the risks of domestic and foreign debt are not completely separable. Financial integration and open capital accounts have encouraged arbitrage between them, and foreigners now own some domestic debt. And a default on that debt, much of which is dollarized or otherwise linked to economic activity, could trigger a crisis for reputational reasons and through the havoc it causes to the banking system and to domestic prosperity. The authors document the growing importance of domestic debt in many middle-income countries, including many with a history of default, and judge that this exposure could lie at the heart of future crises.

THE FEDERAL RESERVE HAS reduced the interest rate on federal funds, its principal policy instrument, by more than 5 percentage points since the end of 2000, yet unemployment continued to drift upward in the first half of 2003. Although the central bank remains dedicated to maintaining economic expansion, interest rates at all maturities are already low, leaving less scope than usual for further monetary stimulus. Especially in such an environment, it is sensible to complement monetary policy with well-targeted fiscal stimulus. But recent developments complicate the prospects for doing so. The federal budget has already undergone a profound reversal, with large deficits replacing the surpluses that a few years ago were projected for 2003 and beyond. And politicians have obscured the distinction between the long-run and short-run budgetary needs of the economy. In the second paper of this volume, Alan Auerbach analyzes the current state of fiscal policy and our understanding of its determinants and effects. He examines how the present budget situation arose, looks for historical evidence on how budgetary and economic conditions have typically affected fiscal policy, assesses what we know about the effects of fiscal policy on the economy, and considers the prospects for budgets at all levels of government in the years ahead.

In January 2001 the Congressional Budget Office (CBO) projected a federal budget surplus for fiscal year 2003 of $359 billion. By March of this year it was projecting a $246 billion deficit. Auerbach provides informative disaggregations of this $605 billion shift. Only about one-sixth comes from cyclical shortfalls in revenue, $126 billion is directly attributed to legislation that reduced taxes, $112 billion comes from higher projected spending excluding debt service, and $42 billion represents higher debt service costs. The remainder, over $200 billion, comes from what the CBO calls "technical factors," a recognition that revenue in 2000 and the years immediately preceding was enhanced by capital gains, income related to options, bonuses, and similar transitory sources of income that, in retrospect, should not have been projected to persist.

The CBO's projections extend ten years into the future and show that the surplus projected for 2011 has declined from $889 billion in the 2001 projection to $231 billion in the March 2003 projection. However, Auerbach shows that the methodology that the CBO is required to follow makes this latest projection particularly unrealistic. It reflects existing tax law, so that most of the tax reductions passed since 2000 are assumed to expire as scheduled, and the fraction of taxpayers paying the alterna-

tive minimum tax (AMT) is assumed to grow steadily. The CBO's methodology also assumes, contrary to recent experience, that discretionary spending will remain constant in real terms as the economy and the population grow. If, instead, expiring tax provisions are extended, discretionary spending remains a constant fraction of GDP, and the AMT is amended so as to keep constant the fraction of taxpayers paying it, the surplus projected for 2011 becomes a deficit of $320 billion. Although that figure differs by well over a trillion dollars from the $889 billion surplus originally projected, such a deficit would not amount to a large share of GDP by recent historical standards. As discussed below, however, Auerbach argues that the long-run budget outlook is considerably more precarious because the CBO projections ignore the looming obligations of the federal government's entitlements programs, primarily those for the elderly.

To get some historical perspective on recent developments, Auerbach turns to an econometric analysis of how the federal government has conducted fiscal policy in the past. He starts with estimates of the full-employment surplus constructed by the CBO, which have the advantage of being available on a quarterly basis over 1955:2 to 2002:4. A regression that explains this surplus by its own lagged value and the lagged output gap (the gap between potential and actual GDP), both as a share of full-employment GDP, shows that fiscal policy over this long period has been countercyclical and has also tended to reduce both budget surpluses and budget deficits. These estimated responses are larger when the sample begins in 1984:3, and largest when it begins in 1993:2, thus covering the Clinton administration and the present administration. To investigate whether the results depend on which party controls Congress and the White House, Auerbach adds to the regression dummy variables for quarters when a Republican was president and when the government was divided between the major parties, as well as variables interacting those dummy variables with the gap and surplus variables. According to these estimates, policy was more responsive to both economic and budgetary conditions during the Clinton administration than it was during the remainder of the sample.

Auerbach notes that changes in the full-employment budget surplus are a flawed proxy for discretionary fiscal actions. One main problem is illustrated by the decline in income tax revenue that shows up as the large technical changes to the CBO projections described above. These revenue

declines reduce the measured full-employment surplus, yet clearly they reflect neither a discretionary change in fiscal policy nor even a development that would be a stimulus to the economy. Timing is another problem. The full-employment surplus changes in response not only to current policy decisions but also to decisions made well before the current quarter; the phased-in tax cuts legislated in 2001 are an example. Finally, the responsiveness of revenue and transfers to economic conditions may vary with time, leading to an inaccurate division of surplus changes into automatic and discretionary components and raising the possibility of a spurious correlation of discretionary changes with the output gap. To deal with these problems, Auerbach constructs an alternative measure of discretionary fiscal actions based on explicit policy changes.

The CBO typically publishes two major revisions of its budget projections each year, the first at the start of the year and the second around midyear. Auerbach uses those changes attributed to legislative action between revisions to derive a roughly semiannual series of changes in revenue and expenditure due to policy. A drawback of this measure is that the continuous series can only be constructed starting in 1984. And the timing and enactment of legislation are again a complication. Some changes are permanent whereas others are temporary, and furthermore, changes legislated in the current period may affect future periods differently. Although he notes these complications, Auerbach first forms a series by simply calculating changes in the current year's surplus that are legislated in the current fiscal year. Regressions over 1984–2003 that explain these changes with the lagged output gap and the lagged surplus show that both have significant, negative coefficients of the same order of magnitude as in the regressions explaining the full-employment surplus over the same period. He next takes account of current legislation that takes effect in future years by adding the discounted value of provisions taking effect up to four years later, using a discount rate of about 50 percent a year, the value that optimizes the fit of the equation. Changes in this more inclusive measure of the surplus are best explained by an equation that uses the level of these projected surpluses in place of the lagged surplus. These same right-hand-side variables also explain changes in the full-employment surplus described above, and with roughly similar coefficients. (The generally countercyclical response of spending, and the response of spending to surpluses, are also supported by regressions based on changes in nondefense discretionary spending.)

Data available at less frequent intervals permit Auerbach to calculate an additional observation for changes legislated in 1981, the first year of Ronald Reagan's presidency, as well as annual changes for 1982 and 1983. When the period is extended to include those observations, the first Reagan year stands out as experiencing an unusually large reduction in the surplus, although it is not the only exceptional year in the sample by this criterion. To examine this episode more fully, Auerbach separates the dependent variable into its revenue and expenditure components and adds a dummy for the change in the 1981 observation in both regressions. These show that the Reagan revolution was a very large outlier both in reducing expenditure and in cutting taxes. Auerbach explores the more general importance of political orientation by dividing up the 1984–2003 period according to whether the president was a Republican or a Democrat, and the results show no important differences.

However, the policy initiatives being legislated in the current fiscal year appear to have long-run budget consequences that are very different from those that the estimated equations would predict. On the assumption that the proposed tax and spending changes are permanent and constant as a fraction of GDP starting in 2004, Auerbach calculates that their cumulative ten-year budget impact would be a revenue loss of $1,455 billion and a spending increase, net of interest, of $725 billion. Under present surplus and economic conditions, the equations predict only a $148 billion revenue loss and $220 billion in higher spending. Thus, viewed in terms of their long-run budget consequences, the present fiscal initiatives stand out as an important ideological initiative, just as the Reagan revolution did. But whereas reduced spending partly offset the effect of the Reagan tax cuts on the surplus, that is not the case this time.

State and local governments face severe financial problems today. Collectively, their current budgets, which measure capital expenditure as the flow of capital services rather than current investment expenditure, were in deficit in 2002 by 0.5 percent of potential GDP; their cash flow budgets, which include investment expenditure, were in deficit by 1.2 percent of potential GDP. To gauge the likely aggregate response of states and localities to these problems, Auerbach investigates how their fiscal actions have responded to budgetary and economic conditions in the past. The analysis largely parallels that of the federal fiscal response. Regressions show that state and local governments typically reduce the sector's full-employment surplus (which is based on their current budgets) by

10 percent of the lagged surplus; there is no significant response to the lagged GDP gap. These results are not surprising in light of the restrictions on deficits that most states face. Auerbach also performs regressions in which the dependent variable is explicit fiscal policy changes rather than changes in the full-employment surplus. For this purpose he uses data on legislated changes in revenue and spending for 1988–2002 assembled by James Poterba from the annual *Fiscal Survey of the States.* These confirm the absence of countercyclical behavior found in the full-employment surplus regressions, and they produce very large estimated effects on revenue and spending from actual budgetary conditions. Separate regressions explaining legislated changes in revenue and spending, using the lagged surplus and the surplus projected before legislated changes as explanatory variables, estimate little effect on current-year revenue but very large effects on current-year spending and next year's revenue. The sum of the effects of lagged and projected surpluses is about –1 on each of these dependent variables. The best-fitting equations add revisions to the projected current surplus from other than legislated changes—the surplus shock—to the regressions. In these regressions the sum of the effects of all three surplus variables on both current spending and next year's revenue remains large but above –1.0. Auerbach suggests that these last estimates are especially vulnerable to simultaneity bias. Using the estimates that omit the surplus shock, he calculates that state and local governments can be expected to cut current spending by $24 billion in 2003 and to increase taxes by $4 billion for 2003 and by $22 billion for 2004.

Turning from the current situation to the longer run, Auerbach suggests that the "real fiscal danger" (a phrase taken from the president's 2004 budget proposal) comes from the future liabilities of old-age entitlement programs—liabilities that are not adequately reflected in current budgetary presentations or in the debates that are informed by them. Auerbach calculates the implicit annual deficits of the Social Security system using the conventional seventy-five-year horizon and the "closed group" concept, which considers the present value of future benefits, net of future contributions, only of age cohorts already participating in the program. The implicit debt of the program calculated in this way rises from $7.7 trillion in 1997 to $11.1 trillion in 2003. (Auerbach adds that the implicit long-term liabilities of Medicare are estimated to be at least as large as these Social Security liabilities.) Because of changing assump-

tions about factors such as interest rates and productivity, which have no clear pattern but on balance reduce the implicit debt, the annual implicit deficits range from $173 billion to $878 billion over this period. Ignoring the effect of changing assumptions, the implicit annual deficit rises from $523 billion to $731 billion and will continue rising as the retirement of the baby boomers approaches. The deficits calculated under either of these concepts, and under reasonable alternative ways of calculating the present value of future commitments, are drastically different from the current cash-flow surpluses of Social Security presented in the budget and in budget projections. How this affects either private agents' behavior or the deliberations of government policymakers is an important question for which there is no agreed-upon answer. However, Auerbach suggests that both could be affected. In particular, he suggests that the much larger total deficits that would result from incorporating the implicit deficits in the budget could, appropriately, restrain government behavior.

Economists have found it difficult to measure the impact of fiscal policy on the economy, and Auerbach reviews a number of reasons why this might be so. Estimation problems arise because the endogenous response of policy to the state of the economy is mingled with the effects of policy on the economy, and separating these requires identifying what was an exogenous shock. Characterizing fiscal shocks poses other problems. Revenue effects are a simple characterization of tax changes, but they do not capture possible responses to changed incentives or the distribution of benefits and losses across income groups. Accounting for expectations is especially difficult. Permanent and temporary changes may have different effects, and whether a fiscal position is seen as sustainable may affect how private agents react to it. Auerbach describes how the Clinton policy program could plausibly have produced an expansionary budget contraction by gradually convincing agents in capital and goods markets that budget discipline would be long lasting. But he notes that this explanation for the late-1990s' boom depends on a complicated set of responses and changes in expectations, making it difficult to verify and leaving many analysts skeptical.

Auerbach notes that the present conjuncture of fiscal policy and the economy is unusual. Interest rates are low, and expectations of future rates embodied in longer maturities have been volatile. State and local governments face unprecedented deficits that will lead to a historic contraction of their aggregate budgets in the immediate future. At the federal

level, official deficits and debt levels are well within the historical range relative to GDP. But the implicit debt and deficits of the entitlements programs are not, and hence in Auerbach's view the United States is in a period of severe fiscal stress. Yet in this situation the president and Congress are pursuing a much larger expansion of federal deficits over the next decade than historical analysis would predict. Auerbach finds it difficult to forecast how the economy will respond to such a policy regime.

SINCE 1995 THE JAPANESE economy has seemed to be in a liquidity trap, a situation last seen in a major industrial economy during the Great Depression. Japanese wholesale prices have been slipping since the end of 2000, consumer prices are lower today than they were in 1998, and growth has been anemic. Monetary policy has appeared powerless to generate a healthy expansion. The overnight interest rate on cash has been within 50 basis points of zero for several years, and virtually at zero most recently, and as a result, the monetary authority has been unable to keep real interest rates from rising as prices fell. Nor has vigorous expansion of the monetary base had an apparent effect on the Japanese economy. More recently, unemployment rates in the United States and most other industrial countries have been rising while inflation has fallen to very low rates, by some measures near zero. This situation, viewed alongside the prolonged deflationary experience of Japan, has heightened interest in the special problems that deflation poses for stabilization policy. Although most economists would agree with Federal Reserve Chairman Alan Greenspan that deflation is not an imminent danger in the United States, they also believe it would pose problems for policy if it took hold. And many would agree with Paul Krugman's emphasis, in a 1998 Brookings Paper, on the need for policymakers to affect private sector expectations in the face of a possible deflation. In the third paper of this volume, Gauti Eggertsson and Michael Woodford present a detailed theoretical analysis of the way policy and expectations interact, providing guidance on the appropriate policy before, during, and after the economy falls into a liquidity trap.

To analyze the implications for the conduct and effectiveness of monetary policy of the zero lower bound on nominal interest rates, the authors utilize an explicitly intertemporal model in which expectations play a central role. The model assumes a representative household that consumes a composite consumption good and supplies each of the various types of

labor that are used in different industries. There is no capital in the model and hence no national saving in equilibrium. The important intertemporal link is the assumption that the household equates the marginal utility of forgone consumption at time t with the discounted expected marginal utility of the increased consumption thereby made available at time $t + 1$. Because of these links, the household's current consumption and labor supply decisions depend on the entire expected future path of short-term real interest rates. Those real rates in turn depend on the nominal rates set by the monetary authority. Real money balances are included in the household utility function, and the demand for such balances can be satiated even at a zero interest rate and with a finite money supply.

To provide a logical basis for sticky prices, the authors adopt the standard assumption that goods are differentiated and produced by monopolistically competitive firms. Prices in a given industry are fixed for a random period of time, and, by symmetry, all firms in an industry set the same price when they have the opportunity. Demands for the differentiated goods are a constant-elasticity function of household income and price. For simplicity, the production function relating output to labor input is the same for each good and subject only to a common exogenous productivity shock. Firms in different industries use different types of labor. Wages, in contrast to prices, are completely flexible, equating the supply and demand of each type of labor.

The authors focus on the optimal conduct of monetary policy and include only a limited role for fiscal policy in their model. Taxes and transfers are nondistorting, and there is no government consumption. However, the government determines the time path of total government liabilities and their composition, which includes non-interest-bearing money. The total is determined by a rule that specifies the real level of liabilities as a function of the preexisting level and current macroeconomic conditions. A similar rule governs the maturity and state-dependent structure of debt. The model assumes Arrow-Debreu markets that are complete across time and states of nature, and it imposes no limits on borrowing. The household optimizes its consumption, taking prices at a given time and in a given state of the economy as given, subject to a single intertemporal budget constraint. Because the government does not consume goods, any saving or dissaving by the government is exactly matched by dissaving or saving by the household.

The authors stress that the entire expected future path of real short-term interest rates matters in their model: changing future monetary actions, which affect future nominal rates and inflation, affect current as well as future output. Although the path of expected short-term rates implies a term structure of rates on assets of longer maturity, the authors show that, given actual and expected short-term rates, open-market operations or government debt policy actions that change the composition of privately held nonmonetary debt have no effect on the economy. This neutrality result arises because the representative household's intertemporal budget constraint is unaffected by such changes. With no government consumption, the household's intertemporal budget constraint mirrors that of the government.

The dependence of the present on the future becomes particularly important when the economy is in a liquidity trap, with the nominal interest rate at zero. In those circumstances changing the quantity of money, or "quantitative easing," is not a separate policy instrument because households are indifferent between holding money and holding short-term government debt. But this does not mean monetary policy is powerless. Even in these circumstances, output and prices do respond to changes in expected future monetary actions; the monetary authority can mitigate the effects of the constraint by affecting the private sector's expectations about the future path of short-term rates. In particular, a credible commitment to lower nominal rates and higher inflation when the economy is no longer in the trap will raise current prices and output. Hence managing expectations about future monetary policy is crucial.

The authors note that various actions that do not directly affect the path of short-term rates may nonetheless be important because they affect the private sector's expectations about future policy. Quantitative easing that expands the money supply today to a level that, if left unchanged, would result in greater inflation in the future may be one way to affect those expectations. Buying long-term bonds when their yields imply higher future short-term rates than the policymaker intends is another. Not only does policy need to be forward looking; when the economy is in a trap, it needs to credibly commit to lower rates and higher inflation in the future than will seem optimal at that later date. As the authors note, this implies that an optimal rule is not just forward looking; it needs to be backward looking as well: to be credible in the future, current policy

needs to honor commitments made previously when the economy was in the trap.

How important is the constraint created by the lower bound on nominal interest rates, and how should policy be adapted to minimize its costs? To address these questions the authors numerically simulate a log-linear approximation of the full model around a zero-inflation steady state. This approximate model can be summarized by two forward-looking equations, one relating current output to expected future output and the difference between the expected real rate and the natural real rate of interest, the other relating inflation to the output gap and expected inflation. Shocks in the full model are captured in the stochastic natural real rate and a "cost-push" disturbance to the inflation rate. The parameters of these two equations are chosen so as to provide plausible dynamics.

The authors first consider a rule targeting inflation. To illustrate how a commitment to inflation can mitigate the effects of the lower bound on interest rates, they simulate the model with an unexpected shock that lowers the natural rate from its steady-state value of 4 percent to –2 percent. Each period following the shock there is a 10 percent chance that the natural rate returns to its steady-state value. If the monetary authority chooses zero inflation as its target, its inability to set a negative nominal rate results in a 14 percent output gap and 11 percent annual deflation while in the trap. The fall in the natural rate reduces output, thereby putting immediate downward pressure on inflation and raising the real interest rate, which in turn further reduces output. These effects are amplified by the fact that there is a 90 percent chance of remaining in the trap with deflation in the following period; the expectation of future deflation increases today's deflation. Setting a target inflation rate of 1 percent substantially reduces the fall in output and prices, because the private sector's expectation of 1 percent inflation once out of the trap stimulates demand. But the rate of deflation is still 4 percent and the output gap just over 7 percent in this case. A policy that credibly commits to 2 percent inflation, cutting the nominal rate to zero when the natural rate falls to –2 percent, avoids any output decline. However, full avoidance is not optimal. The authors note that although commitment to a higher inflation target reduces the output gap in periods when the natural rate of interest falls below zero, it comes with inflationary distortions, both in normal times and when the natural rate is negative. For this reason, the optimal inflation

target for rules of this kind will be positive, but not so high as to avoid all risk of running into the trap.

This simple rule of targeting a fixed inflation rate is not optimal. Because it does not respond to current or historical conditions, it does not optimize the objectives of stabilizing output and inflation. The authors show that one can do better if the policy is state dependent and the policymaker can credibly commit to a history-dependent rule even if, as in their model, past conditions are irrelevant to what can be achieved at present. The authors recognize that it may be difficult to make such policies credible, because a policymaker is committing to actions in the future that will not appear optimal then. But since the best rational expectations equilibrium requires credibility, they first consider optimal policy assuming it is credible before considering how to achieve credibility. They derive the optimal policy response to the shock to the natural rate considered above, assuming the government wants to minimize the discounted expected value of a quadratic function of inflation and the output gap. The optimal interest rate policy has several interesting features. First, it is history dependent. It involves committing the monetary authority to the creation of a future output boom, with accompanying inflation, once the natural rate again turns positive. As in the case of fixed inflation targeting just discussed, this commitment stimulates demand and reduces deflationary pressures when the economy is in the trap. And, as in the case of inflation targeting, the distortionary costs of creating such a boom mean that it will not be optimal to completely offset the fall in the natural real rate. However, the rule improves on fixed inflation targeting by gradually eliminating the boom and the associated inflation following the return of the natural rate to its normal level.

The optimal interest rate policy is a complicated function of the current and past states of the economy as well as expectations of the future, even if one assumes the simplest stochastic process for the natural rate. The authors suggest that a complete description of the state-contingent interest rate rule is unlikely to be feasible and, even if it could be provided, would be extremely difficult to communicate to the private sector. Yet accurate private sector expectations are crucial to the policy's effectiveness. Fortunately, the authors are able to show that the optimal policy is equivalent to a simpler rule that specifies a single target that is a fixed-weight index of the price level and the output gap (which, except for the fact that the

weights are not equal, could be thought of as a target nominal output gap). This target is adjusted up or down depending on whether the actual value of the index is below or above the target, respectively. When in a trap, for example, the index will typically be below the target, with an output gap and deflation. Because of this shortfall, the target is raised. In effect, the policymaker commits to respond to a current shortfall in the index of prices and output by raising the target it will pursue when it is possible to do so.

Because even this simplified characterization of optimal policy is likely to be difficult to communicate to the private sector, the authors also examine the performance of a fixed target rule. They show that such a rule does quite well: it is fully optimal in the absence of the zero bound, and it results in losses only slightly greater than with the fully optimal rule when the zero bound occasionally binds. It may seem surprising that a rule that does not directly take into account the magnitude of a shock to the natural rate does so well. The authors suggest that this reflects the "automatic stabilizer" built into price-level rules. For example, the larger the negative shock to the natural rate, the greater the initial deflation and fall in output and, hence, the greater the implied increase in the target index in the future.

The discussion thus far has concerned the optimal response to an unexpected shock to the natural rate of interest. But suppose such a shock is expected to occur at a future date. Should the policymaker anticipate the shock by changing the current target, or should it leave the target unchanged, or should it even forgo expansionary actions that would be otherwise called for in order to save some scope for more vigorous action in the future—keeping the powder dry? In the author's forward-looking model, the current target level of prices and output should not be changed in response to news that makes a negative shock more likely sometime in the future. However, this does not mean that policy will be unchanged with the news. The private sector's expectations will change with the news, generally leading to lower prices and output at a given interest rate prior to the occurrence of the shock. The policymaker will need to offset this changed private behavior to achieve its fixed target and therefore may encounter the lower bound on rates before the shock hits. The authors recognize that, with more complicated models in which current output depends not just on current and future variables, it may be optimal to raise the target levels of output and inflation before the shock arrives.

Is it possible in this type of rational expectations model for a deflation to be self-fulfilling, so that once the economy enters a trap, deflation and output loss continue forever? The authors show that the answer depends crucially on whether the monetary and fiscal authorities allow the nominal quantity of base money or nominal government liabilities to shrink with the price level, or instead are committed to maintain those stocks at some higher level. If they do not allow shrinkage at the rate of deflation, the self-perpetuating deflationary spiral will not be an equilibrium and will eventually be reversed.

The ability of the policymaker to affect expectations is crucial to the success of policy actions. The authors note that affecting expectations requires more than simply announcing an optimal policy rule, and they discuss policies and actions of the monetary authority that may contribute to the creation of desirable expectations. Publicly announcing and following the optimum price-level rule when there would be little cost to pursuing a different policy—as might be the case prior to encountering the liquidity trap—is one way to make it more credible that the price-level rule will be followed after the event. Once in the trap, adjusting the quantity of money to a level consistent with the target level of prices and output, rather than to their actual level, would communicate the policymaker's target and its intent to meet that target when it becomes feasible. Even if, as in the authors' model, there are no portfolio effects from open-market purchases of long-term bonds, such purchases can signal the authorities' commitment to keep short-term rates low after the natural rate returns to normal. Purchases of foreign exchange are another way the central bank can signal its intentions.

The government can also take actions that give it an incentive to follow the announced policy to inflate once out of the trap. For example, a tax cut, financed by debt or money creation, would provide an additional incentive for the government to inflate, as would government purchases of real assets or foreign exchange. A proposal by Peter Tinsley, that the Federal Reserve issue options to purchase federal funds at a future date on terms consistent with its announced policy, would provide similar incentives. If it issued such options, the Federal Reserve would lose if the funds rate exceeded its previous commitment. The authors conclude that, "given the role of private sector anticipation of a history-dependent policy in realizing a desirable outcome, it is important for central banks to develop effective methods of signaling their policy commitments to the private sector."

DESPITE DECADES OF research, economists' ability to forecast the economy and the effects of policy remains highly uncertain. If anything, there is less agreement today about the theories and models that describe macroeconomic behavior and the potential role of policy than there was forty years ago. Even if economists could agree on the best model of the economy, uncertainty would remain about the size of the parameters that determine how forcefully policymakers should act. In light of these pervasive uncertainties, how should policymakers make decisions, and how can researchers best communicate to them the statistical information they need? In the fourth paper of this issue, William Brock, Steven Durlauf, and Kenneth West recommend a significant change in econometric practice and describe an approach that they believe facilitates intelligent policy formation in the presence of uncertainty.

The authors argue that, in a broad range of contexts, policy evaluation should be conducted on the basis of two factors: the policymaker's preferences (which presumably reflect society's), and the expected distribution of outcomes conditional on policy and available information. This argument implies that the focus of much scholarly research on identifying the true or best model of the economy is of no intrinsic importance; model selection is not a necessary or even a very valuable component of policy evaluation. The authors maintain that model selection is actually inappropriate in many contexts, because it ignores the uncertainty about models that is inevitably present. Instead, the authors advocate the use of some form of model averaging, which provides a formal, disciplined way of proceeding that is not conditional on a particular economic model being true. In principle, their approach is quite simple. The policymaker minimizes his or her expected loss, as derived from a function that expresses the policymaker's preferences over outcomes, where the expectation is taken over the entire set of possible models, parameter variables, and exogenous variables, conditional on the information available and the specified policy. The authors believe that, in this approach, "the observed history of the economy and policy advice are seamlessly integrated," in contrast to conventional approaches, which typically judge a policy's efficacy by the statistical significance of an estimated coefficient from a particular model.

In practice, policymakers are confronted with a variety of possible models, reflecting different theories and specifications of economic relationships, with different implications for the conduct of policy. The

authors believe that uncertainty about the relevance of different models merits special attention, but they defer until later in the paper both that discussion and the discussion of loss functions that incorporate aversion to ambiguity. The authors begin their analysis with the relatively simple case of a specific model for which there is uncertainty about unknowns, θ, and where the policymaker seeks to minimize expected loss. In this simple case, θ represents whatever unknown quantities affect the loss function, including both random shocks that have not been perceived when policy actions are taken and the values of the parameters that determine the effects of policy.

After describing this setup formally, the authors draw out three implications of this approach that differ from conventional econometric practice. First, the relevant uncertainty about θ may not necessarily be described adequately by its expected value and its variance; the entire posterior probability distribution of θ may be relevant. Second, the statistical significance of a policy variable in a regression does not reveal whether or how much that variable should be changed. Whereas discussion of monetary policy rules focuses on loss functions, thereby avoiding this complaint, the economic growth literature is dominated by hypothesis testing and is thus vulnerable to it. Third, the conventional econometric distinction between parameters and estimates of parameters is deficient, as is the practice of reporting standard errors of estimates. More inclusive measures of uncertainty concerning the underlying parameters are needed.

Turning to model uncertainty, the authors observe that it is conceptually straightforward to extend expected loss calculations from one model to several by model averaging, that is, by treating the identity of the true model as an unobserved random variable and weighting the various plausible models by their probability of being the right one. Although it is common econometric practice to report results from some modifications of a baseline model to gauge the robustness of a finding, this informal robustness testing does not provide a way of combining information from these different specifications. The posterior probability of a model depends both on how well it fits the data and on the prior probability that this model is the right one. Although these priors could be based on a variety of sources, including earlier research, a common assumption is that the models being considered have equal prior probabilities. But the authors note that this assumption does not justify placing equal weights

on models in arriving at a policy decision; posterior probabilities also depend on the relative goodness of fit of the various models. They also observe that the uncertainty faced by policymakers is not a simple weighted average of the uncertainty present in each of the alternative models. Differences in the expected outcomes from different models contribute to overall uncertainty and may be much more important than the uncertainty associated with any one of the models.

The authors also consider the possibility that policymakers may treat uncertainty about models differently than they treat other sources of uncertainty. They interpret various experimental results, such as the Ellsberg paradox, as suggesting that agents' behavior in some settings is inconsistent with expected utility maximization, and they note that this can be interpreted as a distaste for model uncertainty. (Ellsberg found that respondents tended to choose a game that they knew they had a 50 percent chance of winning, over a game in which the chance of winning was unknown but had an expected value of 50 percent.) They regard work by Itzhak Gilboa and David Schmeidler and others on ambiguity aversion as providing not just an explanation of behavior but an axiomatic justification for why it may be rational to place more weight on the least-favored model than is warranted by posterior probabilities. The authors propose an approach suggested by Larry Epstein and Tau Wang that formalizes such behavior by specifying a loss function that places weights both on the expected loss computed in Bayesian fashion and on the model that gives the worst outcome. If all of the weight were placed on this second term, policymakers would follow minimax behavior. This can be thought of as the response to an "adversarial agent" who chooses that model which is least favorable to the policymaker. The authors discuss the similarities of this formulation to the extreme bounds analysis proposed by Edward Leamer and the robust optimal control advocated by Lars Hansen and Thomas Sargent. Both of these approaches can also be thought of as embodying ambiguity aversion, but without making use of model probabilities.

The authors go on to a formal analysis showing how a concern for the robustness of a policy choice, meaning a desire to avoid worst outcomes, can affect optimal policy. They begin by considering a core model in which some parameter is fixed, and where the alternative models are "local" in that they differ only in allowing this parameter to lie in a small, well-defined interval around the fixed value. The policymaker's behavior

is assumed to be minimax and is qualitatively different from expected loss minimization. For example, if one were minimizing expected loss, a small amount of uncertainty about the parameter would have no effect on optimal policy. Concern with the worst model causes even a small amount of uncertainty to have a first-order effect.

The authors also discuss other differences that arise. They find that policy responds more or less to some unexpected shock depending on the specifics of the particular model. The authors provide two examples. In the first, the policymaker controls two instruments that have uncertain effects on a single target. In the second, modeled after the analysis of Lars Svensson, the policymaker is concerned about both output, which the policymaker controls, and inflation, which depends on output, shocks, and its own lagged value. In the first example the authors show that, even for a small increase in uncertainty about the variance in the effect of an instrument, robustness considerations will lead to less aggressive use of that instrument and, unless there is an unusual amount of covariance of the influences of the two instruments, to more aggressive use of the other instrument. In the second example they show that greater uncertainty about the feedback of inflation on itself leads to a more aggressive response of policy to unexpected shocks to inflation than otherwise. Robust analysis leads the policymaker to focus on the possibility of higher persistence that amplifies the effect of inflationary shocks unless offset by policy. In this example the authors also discuss the effect of uncertainty about the influence of the controlled variable, output, on inflation. A smaller response is less desirable. But the effect of adding minimax concerns depends on the relative weights in the loss function on inflation and output. The larger the weight on inflation, the more aggressive the policymaker will be in controlling output. Extending this type of analysis to multiple core models is conceptually straightforward, but, as in the simpler examples, the authors show that it is difficult, without knowing the details of the models' specifications, to draw presumptions about how robust analysis affects optimal behavior.

Recognizing that there are many candidate models of the economy complicates the tasks of both researcher and policymaker. The authors propose what they believe to be a useful way to structure the model space. They identify three distinct types of uncertainty: theory uncertainty, which primarily reflects the multiplicity of competing theories and the absence of empirical evidence that would allow one to adjudicate deci-

sively among them; specification uncertainty, which concerns such things as lag lengths, nonlinearities, and the choice of empirical proxies for variables included in the theoretical model; and heterogeneity uncertainty, or uncertainty about the extent to which observations from different settings (for example, countries), different phases of the business cycle, or different historical episodes should be treated as realizations of the same model. They discuss the difficulties that arise if the class of models considered does not include the "true" model, and they suggest that incorporating model uncertainty minimizes the danger from this source.

Although many economists are likely to find the authors' framework conceptually attractive, some will wonder whether it can be put to practical use. The authors provide two examples of how their approach to model uncertainty might be implemented, one involving monetary policy and the other economic growth. Both topics have been subjected to extensive study using more traditional procedures.

In the first example, the authors examine how estimates, using U.S. data, of the performance of a Taylor rule differ when model averaging is used to deal with specification uncertainty about the Phillips and IS curves that drive inflation and the output gap. This example illustrates a frequentist or pseudo-Bayesian approach in which models are given prior probabilities, but the parameters used in each model are not. Theory provides little guidance about the number of lagged values of inflation, the output gap, and interest rates to use on the right-hand side of the equations, especially since they are meant to capture the effects of expectations. To simplify their example, the authors restrict themselves to "backward looking" equations, and, using quarterly data, they allow up to four lags of each of these variables and for the possibility of a structural change in 1984. The sum of the lagged inflation variables in the Phillips curve is always constrained to be 1, and the lagged real interest rate is always included in the IS relation. The various combinations of lags in their specification produce a total of 25,600 models, each of which is estimated by least squares. Expected loss is taken to be a weighted sum of the variances of inflation, the output gap, and changes in interest rates. The specific Taylor rules that the authors use are those that are optimal for different target weights if the economic model is the one estimated by Glenn Rudebusch and Svensson in a 1999 paper. This model corresponds to just one of the authors' models. For each of these Taylor rules, they compute the average loss by weighting the expected loss from each of their models

by its relative likelihood. Somewhat anticlimactically, the authors find that, except for two rules in which the penalty for interest rate variability is low, there is only a small difference between the loss implied by the average of their models and the loss implied by the Rudebusch and Svensson model.

The second application illustrates how a full Bayesian approach can be used to evaluate the effect of tariffs on economic growth. The authors begin with an extended discussion of the existing growth literature and what they see as its deficiencies. First, there are nearly as many theories about growth and as many associated empirical measures as there are countries for which data are available. Individual theories are often tested without recognizing that, typically, one growth theory does not logically preclude another—for example, the theory that political stability affects growth is compatible with any number of other theories. Second, the literature does not deal systematically with heterogeneity uncertainty, for example, whether the coefficient on a policy variable in a cross-country regression can reliably be applied to a given individual country. Third, the literature frequently assumes that statistical significance provides a useful guide to policy evaluation. But using a particular t value as a cutoff not only implies a special loss function, but in practice ignores model uncertainty.

Although the authors regard their illustrative example as no more than a caricature of the actual policy process, since it focuses only on cross-country regressions and ignores the wealth of other information available to help inform policy decisions, in fact it amounts to a substantial study in its own right, dealing with each of their complaints about existing work. The authors include models representing six different categories of theories about the determinants of growth, ranging from theories that focus on exchange rate policies to theories that focus on the characteristics of the political system. Every model includes the variables predicted by the Solow growth model and a variable measuring tariffs on intermediate and capital goods. An effort is made to minimize connections between the theories, so that the probabilities of each theory being relevant are approximately independent. For each theory the authors typically include more than one proxy variable to represent the distinctive growth factor it suggests. They pay special attention to the issue of heterogeneity—in particular, whether the determinants of growth are the same for sub-Saharan Africa as for the rest of the world. They examine the sensitivity of results

to a range of prior probabilities of a difference between countries on this score. Parameter uncertainty is captured by assuming a uniform prior and a Gaussian error distribution. Probabilities are assigned using a tree structure, with probabilities placed on each model, on each of the proxy variables within each model, and finally, for each proxy, on whether or not sub-Saharan Africa differs from the rest of the world. This exercise results in over 8,000 models.

The results are quite interesting. The estimated effects of tariffs on growth are negative and highly significant, and variations between estimates with different priors on heterogeneity are modest. Nor is the model averaging estimate significantly different, quantitatively or statistically, from ordinary least squares estimation of a single model that includes all available policy and structural variables. The authors also report the full range of estimates for their portfolio of models, finding that the tariff variable is always significant and that the response in the model with the least sensitivity of growth to tariffs is roughly 60 percent of that of the model average. Hence the effects of tariffs seem very robust, but the variation across models implies that strong ambiguity aversion on the part of the policymaker would substantially diminish the advantages of tariff reduction in the policymaker's view. Heterogeneity turns out to be less important than the authors expected. Unless the prior on heterogeneity is quite high, the possibility of heterogeneity lowers the mean and increases the posterior distribution of the effects of tariffs only modestly.

The authors use these results to illustrate how statistical significance may not be the best way to think about the policymaker's problem. They calculate the mean and variance of the growth rate for each country in their sample between 1960 and 1985, with and without a 10 percent reduction in tariffs, holding all other regressors constant. Several results stand out. Under model averaging there are large differences in the effect of the tariff reduction across countries. For example, the 10 percent tariff reduction is estimated to raise Botswana's growth over this period by 100 percent but to reduce Burundi's growth by 13 percent. The effects of a change in tariffs on the standard deviation of a country's growth are far smaller than one would guess from looking at the standard deviation of the density of the tariff parameter in isolation. In some cases it actually declines, reflecting the covariance of the effects of tariffs and other growth determinants. This illustrates the importance of computing posterior densities of the outcomes of interest, not focusing on model parame-

ters in isolation. The authors conclude that their analysis supports a policy of tariff reduction for sub-Saharan Africa, unless one has very strong priors that the growth model that applies to the rest of the world does not apply to that region.

The authors view the incorporation of model uncertainty into econometric and policy analysis as still in its infancy. They see a need to extend the analysis in many directions: to account for dynamics and the evolution of the model space over time, to account for nonlinearities and shifts across regimes, and to recognize that policies can and should be updated over time as learning takes place. They also see a host of statistical issues that need attention. But they conclude that "explicit attention to model uncertainty and the use of decision-theoretic methods will prove to be a fruitful direction for future macroeconomic research. At a minimum, explicitly accounting for model uncertainty is an important step in clarifying the limits to which econometric analysis can contribute to policy evaluation."

More than two-thirds of firms in the Standard & Poor's (S&P) 500 currently sponsor defined-benefit (DB) retirement plans, and the value of assets managed by these plans averages around 15 percent of the market value of the sponsoring firm. The defined pension benefits are legal liabilities of the sponsoring firm, and plans must conform to funding requirements specified under federal regulation. Most plans are at least partly invested in equities, so that fluctuations in stock prices affect how much a firm must contribute to the plan in a given year to keep it adequately funded. The boom and bust of the stock market since the mid-1990s has had a significant effect on most DB plans and, thus, on the financial well-being of the sponsoring corporations. But how such gyrations affect these firms' stock prices is an unsettled question; the answer depends on what information is available to shareholders and how they evaluate it. Julia Coronado and Steven Sharpe address this question in a report in this issue.

The authors first describe how firms account for pension funds in their financial reports. Whereas the funding requirements of DB plans are specified by regulation under the Employee Retirement Income Security Act (ERISA) of 1974, guidelines established by the private sector Financial Accounting Standards Board (FASB) govern the accounting requirements for shareholder reports. Many analysts would agree that the net asset

value of a pension fund, with its assets marked to market and its future liabilities appropriately calculated, is a reasonably accurate and transparent measure of the fund's financial position and its contribution to the value of the sponsoring firm. Pension fund net assets could be aggregated with those of the rest of the firm to arrive at a consolidated balance sheet position. Historically, firms have reported pension fund net asset values, but not necessarily in their financial statements and often without sufficient detail. There has also been a lack of uniformity across firms in how future liabilities were valued. In the 1980s, in an attempt to improve accounting for pension funds, FASB imposed requirements (in FAS Statement No. 87) that shifted the emphasis. Under these requirements, firms must report pension fund net asset values in footnotes to their financial statements, and the range of assumptions that firms may use in setting the discount rate for the valuation of future obligations has been narrowed. But a different measure of pension obligations, called the net periodic pension cost (NPPC), must be reflected in the income statement.

The NPPC is calculated as the annual accrued cost of the plan minus the expected return on plan assets. Accrued costs are of two primary types. The first, interest cost, is calculated as the value of pension obligations at the start of the year multiplied by an assumed discount rate (most commonly the rate on triple-A-rated corporate bonds); this measures the increase in the present value of outstanding benefit obligations as employee retirements move one year closer. The second, service cost, is calculated as the present value of *new* obligations, that is, the pension benefits earned by employees during the current year. On the income side, the expected return on plan assets is calculated as a "market related" value of those assets times an assumed rate of return. This assumed rate of return must bear some relation to returns currently being earned as well as to returns expected in the future. For valuing assets, most firms use what amounts to a five-year average of past market values, although FAS 87 permits the use of current market values or some formula in between these extremes.

In contrast to the transparent concept of net asset value, the authors label this accounting for pension funds "opaque." (Although this reflects the authors' view of which measure gives a more timely and meaningful picture of a plan's financial position, it does not imply that net asset values are measured precisely. Market values do exist for plan assets, but the present value of plan liabilities, the other element in the calculation of net

assets or net earnings, depends on assumptions about job turnover and mortality as well as about the appropriate discount rate.) The authors show that the transparent and opaque indicators of pension positions have been weakly correlated in recent years. Between 1993 and 2001 the ratio of net pension assets to net pension earnings swung sharply, from –5.2 at the start to a high of 16.2 in 1997, before falling back to –0.1 in 2001. The divergence widened abruptly between 1999 and 2001. Net assets fell each year with the stock market decline, while net pension earnings rose and then leveled off in step with the moving average of past market values.

Coronado and Sharpe use data for firms in the S&P 500 to address two key questions about how the financial condition of pension plans is reflected in stock values: Does a firm's stock price more nearly reflect the transparent or the opaque measure of the financial condition of its pension plan? And does the stock market value a firm's net pension earnings in the same way that it does earnings from the firm's core activities? The authors use Compustat as their source for historical financial data on individual firms, including information about companies' DB plans, and I/B/E/S International for stock prices and analysts' earning forecasts. Their sample includes all firms that were in the S&P 500 at any time between December 1996 and December 1999, except for a few firms whose pension fund data were incomplete. Their analysis is based on annual data from 1993, the first year for which all the needed pension data are available, through 2001, providing nine years of observations on an average of 490 firms a year. After exclusion of a few firms whose forecasted earnings were negative, the final data set contains 4,359 firm-year observations, 3,335 of which come from firms with active DB plans.

The equation the authors use for estimation is based on the residual income model, according to which a firm's market equity value equals its book equity value plus the present value of its expected abnormal earnings, defined as earnings in excess of the firm's cost of equity capital. Under simplifying assumptions about abnormal earnings, this reduces to a linear equation relating price to book value and earnings per share, with the coefficients on the two explanatory variables reflecting their expected time path—for instance, firms that are expected to grow rapidly have higher coefficients on current earnings. The authors adapt this model for their purposes by separating both a firm's book value and its earnings into a pension fund component and a residual representing the firm's core activity. In making this separation, they treat the accrual of new pension

obligations as (a negative) part of core earnings, so that pension earnings equals the NPPC as defined by FAS 87 less this service cost component. Core book value is the firm's book value less its net pension assets.

If markets respond as predicted by the transparent model, a firm's stock price will be explained by its core book value, its core earnings, and the net assets of the pension fund, all on a per-share basis. Pension earnings per share will not enter. If instead markets respond as predicted by the opaque model, pension earnings per share becomes an explanatory variable, and pension fund net assets will not enter. The authors estimate alternative specifications that allow the data to determine what markets in fact respond to. In addition to the key variables that they test for, all their pooled time-series regressions include time dummy variables and analysts' growth forecasts as well as dummy variables for firms that have DB plans. With pension earnings omitted from the regression, net pension assets has a statistically significant coefficient of 0.76, which is within the range predicted by the transparent model once allowance is made for tax effects. However, when one adds pension earnings, this variable is strongly significant and leaves no role for pension assets in the regression. This finding is robust to a range of specifications and significance tests and thus clearly supports the opaque model. It is also supported in separate cross-sectional regressions for each of the nine years. Only in 1998 does net pension assets replace pension earnings as the significant explanatory variable. Most notably, in 2000 and 2001, when the discrepancies between pension earnings and net pension assets were at their widest, pension earnings dominates the regressions, and the coefficients on pension assets are insignificant (and have the wrong sign).

The other main finding is that the effects of pension earnings on stock price are indistinguishable from the effects of core earnings. Across a range of specifications, the coefficients on both are estimated quite precisely and with no significant difference between them. The authors note that this finding is subject to alternative explanations. Investors may simply look at total earnings and not the composition of earnings. Or they may separately value the contributions of core and pension earnings, and the similar coefficients on the two reflect the average of their sample. To distinguish between these two explanations, the authors divide their sample of firms with DB plans into those with high and low median price-earnings ratios and estimate a fixed effects model on the two groups. They find higher coefficients on both core and pension earnings for the firms

with high price-earnings ratios, indicating that investors simply look at total earnings and do not distinguish between their optimistic view of the firm's prospects and the prospects of its pension fund.

These results are confirmed by regressions based on a dividend discount model as an alternative to the basic stock price model used in the paper. For estimation purposes, the authors modify the dividend discount model to allow for different valuations on core and pension earnings by regressing the logarithm of the ratio of price to core earnings on the logarithm of the ratio of the firm's total earnings to core earnings, its dividend payout rate, and its long-run growth rate, along with the logarithm of firm assets (to capture the effect of firm size) and a dummy variable for whether a firm has a DB plan. The hypothesis that pension and core earnings are valued equally predicts that the logarithm of the ratio of total earnings to core earnings will have a coefficient of 1.0. When this variable is dated one year ahead, the coefficient is indeed not statistically different from 1.0, and under other specifications the authors present, including one with fixed firm effects, the coefficient is above 1.0.

Finally, the authors estimate how much firm values were inflated during the sample period as a consequence of investors' valuing pension earnings rather than pension net assets, the transparent measure that the authors regard as more meaningful. Using the model just described, they first calculate predicted price-earnings ratios for each firm in each year and then calculate the contribution to price made by pension earnings on the assumption that these earnings are valued the same as core earnings. They then estimate each firm's overvaluation as the difference between this contribution from the pension plan to the stock price and the contribution predicted by the transparent model that is based on pension net asset values. By this criterion they find that the average firm with a DB plan is overvalued in each year of their sample, but usually by only a few percentage points. Not surprisingly, however, the overvaluation rises sharply in 2001. In this last year of their sample, the authors estimate that the average firm was 10 percent overvalued and that 10 percent of firms were at least 20 percent overvalued. These overvaluations presumably grew in 2002. But if stock prices continue to recover, they can be expected to switch to undervaluations as the low stock prices of the recent past continue to depress net pension earnings.