Editors' Summary

THE BROOKINGS PANEL ON Economic Activity held its seventy-sixth conference in Washington, D.C., on September 4 and 5, 2003. This issue of *Brookings Papers on Economic Activity* includes the papers and discussions presented at the conference. The first paper examines the Mexican economy since Mexico liberalized its financial markets and concludes that its disappointing performance is due, in important part, to inadequate domestic reforms. The second paper takes a fresh look at understanding the factors behind economic growth over the past forty years, using the experience of eighty-four countries at various stages of development. The third paper looks at how the U.S. productivity surge since the mid-1990s affects our understanding of productivity trends and cycles, and how this might inform future projections. The fourth paper looks for evidence of a bubble in current housing prices.

AFTER MEXICO LIBERALIZED ITS foreign trade and financial transactions beginning in the mid-1980s and joined the North American Free Trade Area (NAFTA) in 1994, it was widely expected that exports would boom, contributing to exceptional economic growth. These expectations have not been met. Mexico experienced a financial crisis in late 1994, from which it recovered with substantial help from the United States. Its exports did rise sharply but then stagnated along with GDP starting in 2001. And on average since liberalization began, Mexico's economic growth has been unexceptional. Mexico's experience, along with the financial crises that beset a wide range of other emerging market economies during the 1990s, raised a number of still-unsettled questions about the benefits and costs of financial liberalization. Although liberalization promoted the inflow of capital from abroad, aiding economic expansion, all too often that foreign capital flowed out just as freely, precipitating crises in some countries and contributing to them in others. In the first paper in this volume, Aaron Tornell, Frank Westermann, and Lorenza Martínez provide a new analysis of how liberalization affects emerging market economies. In their view, liberalization does promote growth by easing financial constraints. But in the process it adds to risk taking in the financial sector, which makes the economy more vulnerable to crisis. How much more risk, and how vulnerable the economy becomes, depend on the ability of lenders to enforce contracts with borrowers and on the prospects for a bailout of the banking sector should a financial crisis occur; these factors in turn depend on certain characteristics of domestic economic and legal institutions. The authors examine these ideas empirically and use the results to inform an extensive analysis of Mexico's recent experience.

The authors begin with an empirical model of the economic linkages among liberalization, financial fragility, and growth. They use a sample of fifty-two countries whose financial markets have achieved a certain level of development, as indicated by a minimum level of activity in their stock markets, and for which data are available for the period 1980–99. The countries are partitioned into seventeen that are judged to have a high degree of contract enforceability (high-enforceability countries, or HECs) and thirty-five that have medium contract enforceability (MECs). The HECs include the Group of Seven large industrial countries and ten others for which an index of the prevalence of the rule of law exceeds a threshold value; all had liberalized both their trade and their financial markets before the start of the sample period. Many of the MECs opened their markets during the sample period, and the authors date these openings by when a trend break occurs in the country's trade or capital flows or when the ratio of trade or capital flows to GDP exceeds a threshold value. The dates identified by this process are similar to those found by earlier researchers.

From these datings and from data on the economic performance of each country, the authors infer a number of stylized facts. One is that trade liberalization has typically preceded financial liberalization in MECs. Another is that both trade liberalization and financial liberalization are associated with faster GDP growth. This conclusion emerges from several specifications of both cross-sectional and panel regressions explaining growth in GDP per capita, all of which include initial GDP per capita, educational attainment (secondary schooling), population growth, and life expectancy as control variables. The estimates indicate that financial liberalization has the greater effect, increasing growth of GDP per capita by between 1.7 and 2.8 percentage points a year in the sample period. A third stylized fact is that financial liberalization is also associated with financial deepening, as revealed by a higher mean level of credit growth, but also with greater financial fragility, captured in negative skewness in credit growth over time. The identification of skewness with financial fragility follows from additional stylized facts about boom-bust cycles in these countries: the typical pattern of rapid credit growth during booms, abrupt declines in credit during the rare crisis, and slow credit growth in the wake of the crisis gives rise to negative outliers in the distribution of credit growth. For the MECs, panel regressions explaining economic growth show a positive association with mean credit growth and with negative skewness of credit growth. However, the sign on negative skewness becomes negative when the trade and financial liberalization variables are added. These results are consistent with the idea that liberalization without crisis is best, that financial fragility is a common by-product of liberalization, and that liberalization is good for growth on balance despite the added fragility.

The authors outline a model that is consistent with these broad results and use it to examine more closely the mechanism through which liberalization affects the economy. (An appendix to the paper provides a formal theoretical extension of the model applied in the text.) Asymmetries between the tradables (T) and the nontradables (N) sectors are a central feature of the model. T-sector firms have relatively unconstrained access to international capital markets through their commercial relations with foreign firms and their ability to pledge export receivables as collateral. Most N-sector firms do not. Except for the largest among them (mainly firms in telecommunications, energy, and finance), N-sector firms depend primarily on domestic bank credit, where their borrowing capacity is constrained because lenders know that contracts are not reliably enforceable. However, the banks are partly protected against the effects of systemic crises by implicit or explicit government guarantees that they will be bailed out in the event of a systemwide crisis. This creates incentives for the banks to channel resources to N-sector firms. These systemic guarantees amount to a subsidy that reduces capital costs and encourages risk taking by both lenders and borrowers. In good times this makes credit readily available to the N-sector as a whole. But the added risk is a threat to the good times. In recent history the risks have mainly taken the form of currency mismatches: banks incur liabilities denominated in foreign currency, while their assets, including loans to N-sector firms, remain denominated in domestic currency. When capital inflows have reversed, leading to a real depreciation, these mismatches have produced bankruptcies and lending crises. A final feature of the authors' model relates the T-sector to the N-sector through real rather than financial effects: an important part of the output of the N-sector consists of inputs to the T-sector. Through this mechanism, not found in most other models that distinguish between tradables and nontradables, prolonged weakness in N-sector output can eventually constrain the output of the T-sector.

By calculating the characteristic behavior of MECs in the three years before and after a crisis and the confidence intervals around the mean behavior, the authors show that their model is consistent with the main features of MEC performance over their 1980–99 data period. A crisis is typically preceded by a real appreciation and a lending boom. During the crisis there is a sharp real depreciation and a banking meltdown, followed closely by a short-lived recession, a sharp decline in credit and investment relative to GDP, and a boom in exports. The ratio of N-sector to T-sector output rises in the years preceding the crisis and then, during the crisis, quickly falls back to earlier levels. The authors associate this pattern with the availability of credit to the two sectors, although they recognize that it is also consistent with the changing relative demand for the output of the two sectors in response to the variation in the real exchange rate. To sort out these two effects, they explain the ratio of N-sector to T-sector output in panel regressions that include as explanatory variables both credit flows and real depreciation as well as the authors' indexes that date financial and trade liberalization and dummy variables for crisis and postcrisis years. In these regressions both credit and the real exchange rate are statistically significant and have sizable effects. The authors also compare Mexico's performance in the three years surrounding its mid-1990s crisis with the characteristic patterns for all MECs. The most noticeable departures in the Mexican case, falling well outside the 95 percent confidence intervals around the mean performance, are rapid export growth after the crisis, a precrisis surge in the credit-to-GDP ratio followed by a sharp decline, and a large and sustained postcrisis decline in the N-to-T output ratio.

Armed with their model and these general results for MECs, the authors turn to the main object of their analysis, a detailed examination of

Mexico's experience since liberalization. They address three main issues: Mexico's mediocre growth performance, the recent stagnation in its exports, and the effects of liberalization and NAFTA. By 1987 Mexico had eliminated most of its nonagricultural trade barriers and had gone from being a very closed economy to one of the most open in the world. Financial liberalization began in 1989 and relaxed or eliminated regulations restricting bank accounts, stock purchases by foreigners, and foreign direct investment. Banks were privatized, and interest rate ceilings and directed lending were eliminated, as were limits on various forms of direct financing by firms. NAFTA, which went into effect in January 1994, greatly reduced the uncertainty confronting investors by codifying the rules of the game that the U.S., Canadian, and Mexican governments would follow on economic and financial matters. However, the authors point to the lack of judicial reform before 2000, when new bankruptcy laws were introduced, as a serious shortcoming in Mexico during this period.

In important respects, Mexico's performance reflected these changes. Between 1985 and 2000, nonoil exports soared from $12 billion to $150 billion, and the share of trade (imports plus exports) in Mexico's GDP rose from 26 percent to 64 percent. The credit-to-GDP ratio rose from 13 percent in 1988 to 49 percent in 1994 (the crisis occurred in December of that year), with credit to the N-sector rising the fastest. GDP growth, which had averaged 2 percent a year in the 1980s, rose to 4 percent during the five years preceding the crisis and was not far from the mean for comparable MECs in the surrounding years. However, performance on some fronts has been unusual and disappointing, particularly in recent years. The N-to-T output ratio declined steadily from 1994 to 2000 and has recovered only slightly since. The real value of credit, which declined sharply with the crisis, continued to fall through 2002, with most of the decline in credit to N-sector firms. And exports and GDP have both stagnated since the start of 2001.

The authors examine these unusual developments more closely. In order to judge how much of the recent stagnation can be attributed to weakness in the U.S. economy, they estimate quarterly bivariate vector autoregressions (VARs) relating Mexican export growth to U.S. import growth and, separately, to growth in U.S. manufacturing. Both VARs show that shocks originating in the United States account for about 40 percent of the forecast error variance, and unexplained changes to

Mexican exports the remaining 60 percent. For the 2001:1–2003:2 period, when Mexican exports stagnated, on average the annual rate of Mexican export growth is underpredicted by around 3 percentage points by the VAR using U.S. imports, and by around 4 percentage points by the VAR using U.S. manufacturing.

In the authors' model, an important factor accounting for this unexplained export weakness is the extended decline in credit to the N-sector and the effect of this credit crunch in retarding N-sector investment, which then causes the N-sector to become a bottleneck to T-sector production, eroding the latter's competitiveness. To support this interpretation, the authors provide a disaggregated analysis of Mexico's export performance since the mid-1990s crisis. From annual surveys of manufacturing, they calculate that N-sector inputs represented on average 12.4 percent of total costs in manufacturing during 1994–99, with this cost share ranging from 5 percent in some food industries to 28 percent for some types of chemicals. Ranking manufacturing industries by this cost share and calculating the cumulative export growth for industries in the top and the bottom 20 percent of this ranking, they show that the exports of the most N-sector-intensive industries fell sharply relative to those of the least N-sector-intensive industries between 1998 and 2002. The authors acknowledge that external factors, such as competing exports from China, have also contributed to the weakness in Mexican exports. But they argue that such external factors cannot explain this strongly asymmetric export response across industries.

The authors also discuss in detail the buildup of resources and lending in the banking system before the crisis, the risks that developed on bank and firm balance sheets as a result of this lending, and the credit crunch that followed. In the early 1990s the large capital inflows that followed financial liberalization, together with a sharp reduction in public sector borrowing, greatly expanded the banks' capacity to lend to the private sector. Because bank liabilities were often denominated in dollars, whereas loans were often in pesos or went to households or N-sector firms whose products were valued in pesos, banks were increasingly exposed to currency risk. In early 1994, expectations apparently changed, because capital inflows reversed, and a few months later a full-blown financial crisis developed. But the more unusual development was the extended credit crunch that followed the crisis. The authors see two main reasons for the credit crunch. The first was a deterioration in contract enforceability:

when many borrowers stopped servicing their debts, and the authorities proved unwilling or unable to enforce payment, a culture of nonpayment developed. This lax enforcement not only led to increased problems for lenders but, the authors believe, also showed up in greater tax evasion and more crime in general. The second reason is the inadequate policy response to the problem of nonperforming loans on banks' balance sheets, which, in the event, kept the share of such loans rising and kept banks from increasing their capital in order to make new loans. Starting in 2000, new steps have been taken on both contract enforcement and bank regulation, but it is too early to judge their final impact.

Using microlevel data that distinguish firms by size, the authors provide evidence on the availability of funds and on the effects it has had on business investment for firms in different parts of the economy. They show that, in 1999, large firms (those with over $2.4 million in fixed assets) accounted for 64 percent of T-sector sales and only 12 percent of N-sector sales. Only large firms are listed on Mexico's stock exchange, and only they issue bonds or equities, and their issuance was sharply higher in the years following the crisis than in the years preceding it. This is evidence that these large firms were insulated from the effects of the credit crunch. Turning to business investment, the authors show that only the top quintile of T-sector firms increased their investment rate between 1994 and 1999. And, in a regression explaining investment rates across the largest firms, they show that cash flow was a more important determinant of investment for nonexporters than for exporters in the wake of the crisis, indicating that N-sector firms were more credit constrained.

The authors conclude with some overall observations about Mexico's performance and the role that liberalization has played in it. They see the benefits of liberalization in the extraordinary growth of exports and foreign direct investment during the 1990s. And they see the lack of structural reform after 1995 and the inadequate and delayed response of policy to the crisis as the main reasons for the disappointing growth of the past few years. From their overall analysis of MECs, they conclude that financial liberalization raises growth rates even though it leads to occasional crises. Although the reforms they discuss for Mexico would reduce the risk of crisis and encourage a prompt recovery should another crisis occur, they judge that liberalization is on balance beneficial even without such reforms. And even though they share the view of most analysts that foreign direct investment is the safest form of capital inflow, they reason that it does not obviate the need for international bank flows, since these provide the main source of external finance for most N-sector firms.

IN THE HALF CENTURY SINCE the classic study by Robert Solow that first attempted to determine the contribution of capital deepening to economic growth in the United States, a vast literature has applied variations of Solow's basic framework to account for growth in a wide variety of countries and time periods. Such studies initially focused on more advanced industrial economies for which adequate data were available, but the development of multicountry data sets in the last decade has made it possible to construct growth accounts for most of the countries in the world. Although these data promise to shed new light on what makes economies grow, comparing economies at widely different stages of development has amplified the importance of a range of measurement issues and has left the answers to central questions about the growth process unsettled. These same data have also been applied to regression analyses of growth, many of which have not been constrained by the production function framework embodied in growth accounting but instead consider different combinations of a broad list of potential explanatory variables. This has led to further proliferation of results but little convergence of expert opinion about the determinants of growth. In the second paper of this volume, Barry Bosworth and Susan Collins address some of the issues that lead to this uncertainty, and they attempt to narrow the range of differences. They do so within a framework that combines the discipline of growth accounting, which establishes the roles of capital formation and other determinants, with regression analysis of a core set of explanatory variables that have consistently been shown to be related to economic growth.

Using data primarily from the World Development Indicators of the World Bank, Bosworth and Collins construct growth accounts for eighty-four countries that together represent 95 percent of world GDP and 84 percent of world population for the period 1960–2000. Their core set of accounts decompose growth in output per worker into growth in capital per worker, increases in education per worker, and a residual, conventionally labeled total factor productivity (TFP). The authors note that TFP captures not just increases in the technological efficiency of factors, as is sometimes assumed, but also myriad other influences on output such as fluctuations in the utilization of capital, political turmoil, changes in government policy, external shocks, and changes in unmeasured factors such as natural resources, as well as any measurement or specification errors. Under the usual assumption that factor input shares are proportional to factor productivities, a factor's contribution to economic growth is its own growth rate times its share. Data on the size of these shares do not exist for many countries. However, where such data do exist, shares appear to change very little over time, and capital's share, once purged of the incomes of the self-employed, appears to be much the same in developing and developed countries. Thus, for all the countries in their data set, the authors assume the same constant-returns-to-scale, Cobb-Douglas production function with a capital share of 0.35, the typical value for Western industrial countries. The authors further assume that capital services are proportional to the capital stock, which is estimated by a perpetual inventory model using investment data from the World Bank that extend back to 1950; labor services are given by the labor force adjusted for educational attainment. For this adjustment they assume that a worker's human capital increases by 7 percent for each additional year of schooling. Data that would allow the analysis to take into account unemployment rates and average hours of work are not available for all the countries in the sample.

Results are given by decade from 1960 to 2000 for each of seven world regions and for the world as a whole. For the entire sample, output grew on average by 4 percent a year, and output per worker by 2.3 percent a year. Of that 2.3 percent, increases in physical capital per worker contributed about 1 percentage point and growth in human capital attributable to education about 0.3 percentage point, leaving 1 percentage point for the residual, TFP. East Asia, even excluding China, was the fastest-growing of the seven regions over these four decades, with output per worker increasing by 3.9 percent a year on average. Interestingly, this rapid growth does not appear to have resulted from these economies "catching up" to the industrial leaders by adopting existing technology; had this been the case, it would have shown up in larger increases in TFP. Rather, increases in physical capital per worker in East Asia were more than twice the global average, and gains in human capital were also above average. Sub-Saharan Africa was the slowest-growing region, with output per worker rising just 0.6 percent a year on average over the entire period. Almost all of this modest growth reflected capital deepening, although this occurred at only about half the rate in the rest of the world. TFP growth in this region was actually negative in three of the four decades as well as over the entire period.

Despite roughly equal contributions of capital accumulation and TFP for the sample as a whole, substantial variation was observed in their relative performance across regions and time. Across regions of the world excluding China, capital's contribution ranges from a high of 2.7 percent a year for the decade of the 1970s in East Asia to –0.1 percent for sub-Saharan Africa in the 1990s. TFP's contribution ranges from a high of 2.6 percent a year in the Middle East in the 1960s to a low of –2.3 percent a year in Latin America in the 1980s.

The authors observe that previous studies, using either regression analysis or growth accounting, have reached surprisingly different conclusions about the relative contributions of capital and TFP. For example, Gregory Mankiw, David Romer, and David Weil conclude that differences in physical and human capital account for roughly 80 percent of the international variation in income per capita, whereas Peter Klenow and Andrés Rodriguez-Clare report that factors captured in the TFP residual account for 90 percent of cross-country variation in growth rates. The use of levels of capital and income in the first of these studies and their growth rates in the second may be one reason to expect differences across these studies and others like them. But beyond that, Bosworth and Collins suggest that different ways of dealing with three key measurement issues are primarily responsible for such dramatic differences.

One of these issues is how to measure growth in capital services. For this some analysts use the investment share of GDP rather than the rate of growth of capital itself. In growth regressions this avoids the need to specify an initial capital stock and assume a rate of its depreciation. But Bosworth and Collins argue that many countries are far from their steady state, and unless the ratio of output to capital is relatively constant, the investment rate can be a poor proxy for growth in capital. The authors find it implausible that this ratio is constant across a diverse sample of countries, and so they argue for using direct measures of capital growth rather than the investment share. And they show that, in their own broad sample, there is very little correlation between the change in the capital stock and the mean investment rate, even over a period as long as forty years. The newly industrializing economies of Asia, for example, have very high capital growth rates but do not devote unusually large shares of output to investment, whereas Guyana and Zambia, with conspicuously slow growth in output and of capital, have relatively high average investment shares. As would be expected given these facts, in a wide variety of growth regressions that differ in terms of the time period examined and the menu of explanatory variables, the authors find that changes in the capital stock explain far more of growth in output per worker than investment rates do.

A second measurement issue is what price concept to use. Bosworth and Collins argue that the choice of international prices (that is, prices adjusted for purchasing power parity), rather than prices in local currency converted at current exchange rates, makes sense in cross-country comparisons of standards of living but is less obviously suitable for growth accounting. In particular, in a growth accounting context, the authors believe that local currency prices should be used, because firms base their production decisions on the relative prices of capital and labor in their domestic markets. The choice makes a substantial difference in the results. Conversion to international prices can dramatically change measured investment shares, systematically reducing them in low-income countries while raising them in most developed countries. The authors report that, in their eighty-four-country sample, the correlation between average investment measured in these two ways is only 0.52. They also find that, across countries, there is almost no correlation between the investment share measured in domestic prices and income per capita, whereas when international prices are used, a strong positive correlation is observed. The importance of the choice of pricing metric is therefore central when trying to explain differences in growth in output by investment's share. The choice also makes a difference in explaining differences in the level of income. It is much less important in equations explaining growth in output by the rate of growth of capital, since the ratio of international to domestic prices is relatively constant.

The third issue discussed by the authors is conceptual. Investment in capital is potentially endogenous: it may well respond to changes in TFP. In a growth regression this creates an obvious econometric problem, biasing upward the estimated coefficient on the growth of capital. But even if the econometric problem is absent, as in the authors' growth accounting where the elasticity of output with respect to capital is assumed to be known, the question of whether capital should get "credit" when the output change was caused by a change in TFP remains. Some researchers limit capital's contribution to that resulting from increases in the capital-output ratio, implicitly giving credit to TFP, or to any other factors recognized in the growth accounting (such as human capital), for the capital additions needed to maintain the capital-output ratio. Some go further and also credit TFP for the changes in human capital required to keep up with output. Such assumptions automatically increase the credit given to TFP. The authors note that these assumptions give extreme estimates. Even if one wants to attribute to technical change whatever growth in output arises from increases in capital that that technical change induces, many unmeasured factors other than technical change will be captured by the TFP residual and may affect both output and capital. Furthermore, causation can run two ways, with technical change itself induced by investment embodying new technology. The authors note that in most growth accounting studies, as in their paper, the focus is not on the deep, ultimate causes of the changes in observed factors but on the proximate sources of growth. For that purpose the growth in capital, rather than in the capital-output ratio, is the relevant variable.

In the authors' growth accounting, what fractions of the variation in growth in output per worker across countries are explained by changes in the measured factors, capital and education, and what fraction by the TFP residual? The authors report a decomposition of the variance of growth using each country's growth over the period 1960–2000 as an observation. If the growth rates of factors and of TFP were uncorrelated, the variance of output growth per worker would be the simple sum of the variances of the growth attributed to each of the components. But if they are not, as in their sample, the calculation also involves the sum of the covariances. The authors report the decomposition of output per worker for various assumptions about the allocation of these covariance terms. If the covariance for any two components is split equally between them, growth in the capital-labor ratio explains 43 percent of the cross-country variance and TFP growth explains 54 percent. The contribution of growth in the capital-labor ratio ranges from 27 to 57 percent, depending on whether it is credited with none or all of its covariance, respectively, with education and TFP. Growth in education explains only 3 percent when the covariance is divided equally. Weighting countries by their population modestly reduces the estimated contribution of capital and increases the contribution of education. Not surprisingly, if growth in the capital-output ratio is used as the measure of capital's contribution, that contribution falls to 12 percent and TFP's contribution rises to 83 percent with equal division of the covariance. The authors note that the covariance between capital's contribution and TFP switches from positive, when growth in the capital-labor ratio is used for capital's contribution, to negative, when growth in the capital-output ratio is used. This suggests that even for those who want to credit TFP for induced changes in capital, using the capital-output ratio is going too far. Bosworth and Collins conclude that both capital (physical and human) accumulation and improvements in economic efficiency are central to the growth process, and that, for most purposes, focusing on which is more important is misplaced.

It is a widely held belief that human capital is important to productivity and therefore that investment in education is an important element in any strategy for development. This view is buttressed by a large body of microeconomic evidence that finds a strong relationship between education and earnings. But many studies at the macroeconomic level have been unable to find a correlation between economic growth and increased educational attainment of the population. In the growth accounting just discussed, education was a relatively minor factor in explaining differences in growth rates. Bosworth and Collins survey a number of micro- and macroeconomic studies, discuss possible reasons for their divergent results, and explore the sensitivity of the macroeconomic results to the use of different measures of human capital and specifications of growth equations. They identify three major reasons why the micro- and macroeconomic studies may come to different conclusions. One is that microeconomic studies typically estimate private, not social, returns, and so their estimated returns need not show up in GDP. It is sometimes argued that an important role played by education—perhaps as important as whatever education adds to an individual's human capital—is simply identifying those individuals who have high native ability. Although such labeling may improve the allocation of labor among tasks of differing requirements, and to that extent has some social value, such sorting is likely to have far greater private than social returns. It is difficult to design studies that can distinguish between these types of return. But the authors note that the natural experiment provided by instances of compulsory education shows little evidence that private returns are biased upward. Hence they do not see the distinction between private and social returns as resolving the puzzling difference between the macro- and the microeconomic results.

Researchers recognize that measurement error, particularly in samples that include extremely diverse economies, may account for the lack of association between economic growth and education in econometric studies. Using measures of education that correct for classification changes in censuses of some industrial countries, some have found substantial effects across that relatively homogeneous set of countries. Others have found particular indicators of education that yield significant results for a broader list of countries. Bosworth and Collins analyze some of these data sets but find that although they are likely to show high correlations in levels, they have much lower correlations in changes, particularly over decades, and so can give quite different results. They explore various strategies for dealing with these measurement problems, including the use of instrumental variables and the use of a composite measure for education, but conclude that although there is substantial evidence of measurement error, there is no convincing way to choose among the different measures.

A third problem in macroeconomic studies of the contribution of education to growth is the difficulty of measuring differences in educational quality. This is obviously a much greater problem for international comparisons than for studies of education within a country. The authors report that an educational quality index generated by Eric Hanuschek and Dennis Kimko for a substantial number of countries, using indicators such as enrollment rates and expenditure per student as well as years of schooling, population, and regional dummies, had a strong correlation with changes in GDP per capita. Bosworth and Collins reestimate and extend the Hanuschek-Kimko measure to their full set of eighty-four countries. They perform a number of growth regressions, always including growth in capital per worker, one or more educational variables (initial years of schooling, a simple average of estimates of human capital from two independent studies, and the index of educational quality), and in some cases other characteristics of the economy in 1960, such as income, life expectancy, and institutional quality. They find only weak evidence of the contribution of education to growth. Growth in human capital is significant only if the coefficient on capital per worker is constrained to its hypothesized value of 0.35, and the significance of the educational quality index disappears if institutional quality is included.

The authors turn next to a more detailed regression analysis relating countries' growth over the 1960–2000 period to some basic measures of initial conditions, external shocks, and policy. The specification is largely drawn from the existing empirical literature. They then use that specification to explore the extent to which the various determinants operate through capital accumulation and through TFP. In their basic regression the authors include five variables they consider exogenous—the initial values of income per capita, life expectancy, and the logarithm of population, along with a trade instrument (as a measure of a country's predisposition to trade) and a geographical factor that is a composite of average days of frost and the proportion of the country's area that lies in the tropics. They experiment with proxies for institutional quality and settle on a measure constructed by Stephen Knack and Philip Keefer from information in the *International Country Risk Guide.* This variable is not predetermined, since it is based on 1982 data and may well be endogenous, reflecting actual growth and policy. These variables by themselves explain 75 percent of the variance in growth of output per worker over the forty-year period, and all are statistically significant. The coefficient on initial income is negative and highly significant, implying convergence. The coefficients on the other variables are all positive, implying that initial life expectancy, trade, and institutional quality all have favorable effects on growth. Adding three policy indicators—the average rate of inflation, the government budget balance, and a measure of openness to trade—improves the fit of the equation only slightly, despite the likelihood that these variables are endogenous. The variables have the expected sign, but only the budget balance is statistically significant. The authors find it striking that there is relatively minor evidence of a direct role for conventional government policies in improving growth performance.

Since actual growth over the period is the sum of the contributions of growth in capital per worker and the TFP residual calculated in the growth accounting, separately regressing these two on the same list of variables allocates the influence of each on growth between a capital and a TFP "channel." The results are suggestive. Institutional quality has a more significant influence on TFP, and trade a more significant influence on the contribution of capital. The other variables have significant effects on both. Adding regional dummies has a modest effect on these coefficient estimates but typically reduces their statistical significance.

Regressions for growth using data over a forty-year period necessarily obscure the remarkable collapse of growth in much of the global economy in the last two decades. Bosworth and Collins document this decline. For their sample, growth in output per worker slowed from an average annual rate of 2.5 percent in 1960–80 to only 0.8 percent in 1980–2000. They show that the decline was similar in industrial and developing countries and is apparent in all regions except South Asia. Their growth accounting shows that lower rates of growth of both physical capital and TFP were important contributors to the widespread decline. India and China are dramatic exceptions. China's growth rate increased by almost 5 percentage points between the two subperiods, and India's by more than 2 percentage points. Since these two countries represent fully 45 percent of the population of the authors' sample, the average growth rate actually increases between subperiods when country observations are weighted by population, giving quite a different picture of what the "typical" world citizen experienced. Weighting observations by GDP results in a decline from the first to the second subperiod of 0.9 percentage point, an intermediate result.

In order to understand the reasons for the general decline in growth, the authors perform separate regressions for the two subperiods, using the same specification that was used for the entire forty-year period. The values of six of the nine variables change between subperiods, but the other three—institutional quality, the trade instrument, and geography—are fixed. Overall, the regressions for the entire period and those for the two subperiods are quite similar. Budget policy appears less important in the second subperiod, whereas geography, institutional quality, trade, and the measure of trade openness all appear more important, with higher statistical significance. However, changes in the variables between subperiods do not explain the general decline, accounting for only 15 percent of the actual change in GDP growth between subperiods. Changes in the variables imply faster growth in the second subperiod for all regions and for developing countries as a group; they imply slower growth for industrial countries but account for only about half of the actual slowdown. It is clear that most of the decline in the average growth rate is reflected in the shift in the intercept and changes in the coefficients on geography and trade. Formal tests show that the changes in structure between the two subperiods are statistically significant. The subperiod regressions allocating the effects of the explanatory variables between capital deepening and TFP show that changes in the effect of geography, institutional quality, and trade are concentrated in the TFP component.

The authors believe that the primary reason for the lack of success in explaining the growth decline is that the regression analysis focused on identifying factors that explain performance across countries rather than factors that would explain variations over time. They note, however, that the equations do predict an acceleration of growth for China, largely due to a large improvement in life expectancy, and predict a slowing of growth in the industrial countries, largely due to a deterioration in government budgets and a reduced role for convergence within this group. The authors also examine growth regressions for various subgroups of countries, finding surprisingly small differences between determinants of growth between higher- and lower-income countries over the entire time period, but also evidence that the shifts in parameter estimates across periods are more substantial for the lower-income countries.

The authors conclude that growth accounting and growth regressions can yield consistent and useful results. They believe their study confirms that a very large portion of cross-country variation in economic growth over the past forty years can be related to differences in initial conditions and government institutions, and that there is robust evidence of convergence. However, they find that the difficulties of accurately measuring educational quality and attainment leave unanswered the contribution of education to growth. And they are disappointed that the variables shown to be important for explaining differences in growth across countries provide little insight into the reduction in average growth rates during the last twenty years, and that they find only a minor role for economic policy in boosting countries' growth performance.

EVER SINCE PRODUCTIVITY accelerated in the second half of the 1990s and unemployment fell to its lowest level in thirty years, analysts have worked at understanding the reasons behind this surprising performance of the economy and the implications for the future. Because productivity has a pronounced procyclical component, it had been unclear what portion of these rapid productivity gains was a transient cyclical development. Thus the renewed acceleration of productivity in the past several quarters, resulting in rapid average gains during a period when unemployment rose, has been informative, suggeesting that there is indeed an important continuing component of the productivity resurgence. In the third paper of this volume, Robert Gordon analyzes these developments and applies his analysis to three important issues: interpreting the rapid productivity gains of recent years, understanding the behavior of productivity over a business cycle, and projecting the likely trends in productivity and output growth in the next twenty years.

Estimating the underlying trend in productivity is the starting point for Gordon's analysis. Total output of the economy, GDP, can be decomposed into the product of five terms: output per hour (productivity), hours per employee (average hours), employment as a fraction of the labor force (the employment rate), the labor force as a fraction of the working-age population (the labor force participation rate), and the working-age population. Because productivity is reported quarterly for the nonfarm business sector, it is useful to work with productivity and average hours for that sector and then get back to GDP by adding two more terms that relate total output to nonfarm business output (what Gordon calls the "mix effect"), and nonfarm business employment (from the Bureau of Labor Statistics' payroll data) to total employment (from the BLS's household survey). Since the product of these seven terms is GDP, the growth rate of GDP is given by the sum of their growth rates.

To illustrate the contribution of each term to GDP growth over time, Gordon divides the period 1954-2001 into six intervals between benchmark quarters that he regards as roughly comparable points in business cycles. Productivity growth makes a positive contribution to GDP growth in all six periods, but its average annual growth rate varies between 1.2 and 2.7 percent across periods. Population growth contributes between 1.1 and 1.9 percent over the intervals, and a rising participation rate contributes between a barely positive amount and 0.7 percent. Average hours decline between zero and 0.6 percent. The other terms are smaller and vary between positive and negative over the intervals.

To go beyond this rough characterization using trend benchmark quarters (what he calls the trends-through-benchmark, or TTB method), Gordon turns to a more formal analysis of productivity using filter techniques. One is the familiar Hodrick-Prescott (H-P) filter, which estimates a time-varying trend by simply smoothing time series on actual productivity. Gordon experiments with a number of alternative smoothing parameters and settles on one by judgmentally examining the variations in trend it produces. Because cyclical variations in productivity are so pronounced, the estimated H-P trends are contaminated by an unknown amount of cyclical variation. That leads him to try a second technique, a Kalman filter that estimates a stationary cyclical adjustment along with a time-varying trend. For his cyclical adjustment, Gordon utilizes the output gap calculated as the difference between actual GDP and the broken trends using his TTB method, and he uses the current and four leading values of changes in this gap as additional explanatory variables in the Kalman filter. He chooses a degree of smoothing that, with the cyclical adjustment suppressed, gives an estimated trend resembling the preferred H-P estimated trend. The final Kalman filter trend is then obtained by reestimating using this degree of smoothing, but including the output gap variables. The resulting trend is noticeably different from the H-P estimate and relatively free of transient cyclical variation.

The differences in the estimates of trend growth rates using these two methods are most noticeable in the early years of the estimation period, 1955–70, and at the end of the sample, 2000–03. The H-P filter estimates show a strong surge in productivity from 1955 to 1963, followed by a steady decline to the late 1970s. The Kalman filter estimates show little variation in productivity growth until the early 1970s, followed by a steady decline to the late 1970s. The terminal growth rate of the productivity trend in mid-2003 is estimated at 3.4 percent a year using the Kalman filter, 0.7 percentage point faster than the H-P estimate. Confronted with these differences, Gordon chooses an average of the productivity trends estimated by these two techniques as his final estimate. He then repeats the procedure to estimate time-varying trends for five of the remaining six terms in the GDP identity. (Where trends are estimated from growth rates, the levels of the trend are calculated by constraining the average ratio of actual to trend to equal zero over the entire 1954:4–2003:2 sample period.) The exception is the employment rate, which is 1 minus the trend unemployment rate. Gordon identifies trend unemployment with the NAIRU. Because filter estimates of the employment rate trend do not imply a very slowly changing NAIRU like that he has found successful in previous work, Gordon relies instead on the TTB method. This results in an employment rate trend that rises gradually from 94 percent in the late 1980s to 95 percent by 2001 and identifies the period from late 1996 to late 2001 as an episode in which unemployment remained below the NAIRU. From the trends in each of the components, Gordon calculates the trend values and growth rates for GDP.

Gordon uses these estimates of trends in GDP and its components to examine the cyclical behavior of the economy as originally expressed in

Okun's Law. The log ratio of actual to trend values measures the percentage deviation from trend. For each of the components, Gordon regresses the change in this deviation on the change in the deviation of GDP (with zero to four lags, except in regressions for productivity, where four leads are used), the lagged dependent variable (with four lags), an error correction term (the lagged log ratio of actual to trend), and an end-of expansion (EOE) dummy variable that uses seven leads, which his earlier work found to be important in explaining cyclical productivity. Confirming Okun's original analysis, he finds a significant relationship between deviations in GDP and deviations in productivity, average hours, and the employment rate, and smaller but also significant effects of the variables relating nonfarm business output to GDP. (In the 1960s Arthur Okun also found effects from the labor force participation rate, but these are not significant in Gordon's analysis when the last forty years are added to the data.) The estimated short-run effect of a 1 percent deviation in output from its trend is a 0.50 percent deviation of both productivity and the employment rate and a 0.29 percent deviation in average hours worked. The other terms, taken together, have small coefficients. In steady state, the productivity deviation falls to 0.33 percent and the employment rate deviation to 0.41 percent, and the average hours deviation changes only slightly. The steady-state relation between GDP and the employment rate implies an Okun's Law of 2.4.

The surge in productivity in the last two or three years poses a special challenge to any analysis. Gordon discusses this period in detail and considers the predictions of alternative specifications of his productivity equations in attempting to shed light on this episode. Although the productivity trend he uses, which averages the H-P and Kalman filter estimates, is rising by just over 3 percent a year during this period, the equations still consistently underpredict actual productivity growth in the quarters starting in early 2000. Gordon considers whether the typical EOE effect may have been absent in the 2001 recession and the current recovery. His estimates omitting the EOE effect do track actual productivity in 2000 and early 2001 closely, but they generate much larger underpredictions in the subsequent quarters. The equations also substantially overpredict productivity growth during 1993 to 1999, indicating the difficulty in quantifying the changing trend throughout this entire period. Given the decline of GDP relative to its trend from 2000:2 to 2003:2, the declines in both the employment rate and average hours were near their predicted values, although the participation rate declined by 0.26 percent a year more than predicted. Thus developments in the labor market were not unusual given the weakness in the economy. However, compared with two earlier periods that Gordon examines, unusually large prediction errors emerge in the two terms that reconcile GDP and nonfarm business output. The ratio of GDP to nonfarm business output declined slightly, whereas it was predicted to rise by 0.58 percent a year, and the ratio of payroll to household employment fell by 0.13 percent a year more than predicted.

The acceleration of productivity during the past decade has been analyzed extensively, and Gordon critically reviews a number of possible explanations that others have put forward. One of these focuses on extreme cost cutting by firms as a factor behind the 2002–03 surge in productivity. Gordon agrees that a number of developments put severe pressure on corporate managers to cut costs, which led to unusually large employment cuts. Reported profits declined sharply after 2000, in part because of the unwinding of accounting tricks that William Nordhaus identified in a recent Brookings Paper, and the stock market collapsed. Compensation that was geared to stock options vanished, and many pension funds that were invested in equities became seriously underfunded. Another explanation centers on the delay that user industries often experience before realizing the benefits of their investments in new technology hardware and software. This explanation was originally articulated by Paul David, who drew an analogy between the delayed benefits that followed the introduction and spread of electricity roughly a century ago and those that followed the introduction and spread of computers and related technologies in the last few decades. This explanation sees the recent surge in productivity in user industries as the benefit of investments in information and communications technology (ICT) made in the 1990s. The delay arises from the need to reorganize business practices in order to achieve the benefits of the new investments.

A third and related explanation is based on the idea that measured ICT investment is accompanied by unmeasured investment in intangible capital. This unmeasured investment consists of the resources that firms devote to reorganizing, retraining, and otherwise acquiring the human capital that is the necessary complement to the measured ICT investment.

Because the use of firms' resources for unmeasured investment is not captured in measured output at the time they occur, the intangible capital hypothesis asserts that output and productivity growth were understated in the late 1990s when ICT investment boomed, and that that understatement was partly reversed in recent years. Gordon estimates that the effects of such mismeasurement, even if present, could account for only part of the surprising recent surge in productivity. He also discusses the possible role of greater outsourcing, labor market flexibility, and mismeasurement of hours worked but finds the evidence for effects on productivity from these sources to be slight or ambiguous.

The trend growth in productivity is the principal determinant of growth in real wages and, together with growth in the other elements of the GDP decomposition analyzed above, determines the growth of potential GDP. Projections of GDP growth, in turn, are an essential ingredient of long-run budget forecasts and forecasts of future revenue for Social Security and Medicare. Gordon applies the results of his historical analysis to project potential GDP growth over the next twenty years and compares these projections with the current official projections for Social Security. He offers several reasons for not projecting his 3.05 percent estimate of the current trend in productivity growth for this long a period, noting that productivity growth over the past twenty years averaged just 2.03 percent and that estimates of the trend changed sharply during that period. One is that the extreme cost-cutting efforts that may have contributed to the recent productivity surge cannot continue indefinitely. A second is that the overstatement of recent growth implied by the intangible capital hypothesis will soon end. A third is that the pace of fruitful applications of ICT innovations is likely to slow. And a fourth is that educational attainment in the United States is reaching a plateau, so that the contribution of improving labor quality to productivity will probably be smaller in the future than it has been in recent decades. Taking all this into account, Gordon judges that productivity growth over the next twenty years is likely to average between 2.2 and 2.8 percent a year, and he uses 2.5 percent for his central projection of productivity growth over this period.

Turning to the other elements in the projection of potential GDP, Gordon takes issue with the population growth projections currently adopted by the Social Security Administration's Technical Panel on Assumptions and Methods (TPAM). Population growth depends on fertility rates, mortality rates, and net immigration. Gordon notes that the U.S. total fertility rate in recent years has been very near the 2.1 births per female of child-bearing age that would, absent immigration, maintain a constant population in the long run. Although this is well above the fertility rates of 1.3 to 1.6 that characterize most other advanced industrial countries, Gordon agrees with most demographers that the U.S. rate will remain near 2.0, and he discusses some of the differences between the United States and other rich countries that support this projection. He offers no disagreement with the TPAM's projection that mortality rates will decline in the future at the same 0.84 percent a year rate that has characterized the last fifty years, and he accepts its projection of no change in average hours per household employee. However, he argues that the TPAM projection of net immigration is too low. Net annual immigration ranged between 0.6 and 1.0 percent of the population in the early years of the twentieth century, declined to near zero between 1933 and 1947, and gradually rose to 0.4 percent in 2002, when net immigration (including illegal immigrants) was 1.4 million. The TPAM projects a decline to a steady state of 900,000 a year after 2020. Gordon notes that net immigration has risen an average of 3.75 percent a year since 1970, and he questions whether that growth will switch to an absolute decline in the years ahead. He discusses alternatives to the TPAM projections that he finds more plausible. These make a considerable difference in the seventy-five-year projections made for Social Security's purposes, but only a small difference over the next twenty years. For that period, and taking account of all three elements of population growth, Gordon projects that the working-age population will grow at an average rate of just under 1 percent a year, somewhat slower than the 1.2 percent annual growth over 1987–2003, but well above the TPAM projection for labor force growth of 0.63 percent a year. He also assumes no change in the trend employment rate or the labor force participation rate, and he assumes a continuation of the recent 0.2 percent annual decline in nonfarm output relative to GDP. Taken together, these projections lead to projected growth of potential GDP averaging 3.28 percent a year over the next twenty years.

Although Gordon does not attempt to model uncertainty in this projection in any formal way, or to quantify the uncertainties in the individual elements from which it was generated, he is well aware of them. Growth in productivity, population, and hours per employee are the three components of the output identity that he regards as the most uncertain for long-term projections. His own analysis of productivity identifies substantial variations in its growth trend over the past half-century and a doubling of the estimated trend over just the past ten years. Economists have little understanding of the forces behind these variations and, Gordon notes, disagree about whether projections for the next two decades are best informed by average growth over the past two years, the past eight, or the past twenty. He concludes that "the only safe forecast is that we will be surprised sooner rather than later."

SINCE 1995, HOUSING prices in almost every U.S. metropolitan area have been rising far more rapidly than incomes or than prices in general. And despite the recession in 2001, prices of single-family homes, the volume of house sales, and the number of housing starts have all remained at near-record levels. Indeed, the resilience of housing sales (and of automobile sales) has been a major factor preventing a deeper recession. In many respects the rapid run-up of home prices, which is mirrored in most other advanced economies, resembles the housing boom of the late 1980s, which was most dramatic in California and Massachusetts. The fact that those earlier housing price surges in Los Angeles, San Francisco, and Boston were followed by housing busts in those areas has many observers concerned that the United States today may be in a housing bubble, with the risk of a damaging collapse. In the fourth paper of this volume, Karl Case and Robert Shiller examine a range of evidence bearing on the likelihood of such an outcome.

The authors distinguish between a housing boom and a housing bubble but recognize that it is difficult to tell the two apart. A housing boom occurs when rapidly rising housing prices and high rates of construction reflect favorable fundamentals, such as falling interest rates or rapidly rising incomes or wealth. A bubble, in contrast, is a situation in which such price increases reflect expectations about future price increases that are not justified by changes in fundamentals. Although such expectations can be self-fulfilling for a time, as expectations of price increases beget price increases, they are inherently unstable: a shock that disrupts those expectations can induce a collapse.

Case and Shiller first examine the extent to which changes in fundamentals can explain observed fluctuations in home prices. They construct state-by-state home price data for the period 1985:1 to 2002:3, using repeat-sales and appraisal indices, scaled to median home values in 1999.

They find that the variability of home prices over time differs widely across states. For the seven states with the least volatile prices, the average ratio of home prices to personal income per capita is relatively low, averaging about 2.3, and the fluctuations in this ratio are quite modest, with a standard deviation ranging from 0.04 to 0.09. In simple regressions of home prices on income per capita, the latter by itself explains over 99 percent of the variance of home prices for four of these states, and 96 percent or more for the other three. Although this leaves little for other variables to explain, the authors do find that including a number of other fundamentals—changes in population and employment, the unemployment rate, housing starts, the mortgage rate, and mortgage payment per $1,000 of loan value—raises the R^2 to 99 percent for all the least price-volatile states.

The picture is quite different for the eight states where home prices are the most volatile (which include California and Massachusetts). For these states the ratio of home price to income, averaged across states and time, is almost triple that of the seven states with the least volatility. The standard deviations of this ratio range from 0.52 to 1.34, averaging over ten times the average standard deviations of the low-volatility states. Not surprisingly, for these states income per capita by itself explains much less of the volatility of home prices, and other fundamentals appear more important. In a regression in which income per capita is the only independent variable, the unexplained variation in home prices, which was at most 4 percent for the least price-volatile states, ranges from 17 to 55 percent for the most price-volatile states. Inclusion of the other fundamental variables listed above typically cuts the unexplained variation roughly in half, which leaves ample room for some of the price variation to be explained by a bubble.

The authors examine the explanatory potential of the fundamentals in a variety of other regressions, using both levels and rates of change in variables, with similar results. They recognize that the choice of functional form and of the explanatory variables in home price equations is somewhat arbitrary. And they stress that some of the variables in these regressions are endogenous, cautioning against taking them as strong evidence of the importance of fundamentals. In particular, mortgage rates may be low when home prices are weak, because monetary policy is normally expansive when economic conditions, including housing demand, are weak. Hence the effect of mortgage rates on the demand for housing may be understated. Similarly, housing starts may reflect shifts in supply and therefore have a negative, but underestimated, effect on prices; but they also respond to shifts in demand, biasing the coefficients upward and possibly changing their sign. Nonetheless, it is noteworthy that, in all the most price-volatile states except Hawaii, actual home prices greatly exceed those forecast by the equations for 2000–02, in several cases by more than 25 percent. The authors conclude that their analysis of fundamentals does not reject the hypothesis that a bubble exists in these states.

Households' expectations about future housing prices play the dominant role in explaining bubbles. Economists are often skeptical of using what agents say about their beliefs and behavior, as opposed to prices and quantities in actual transactions, as evidence to test an economic hypothesis. But Case and Shiller believe that carefully constructed surveys, administered close to the time when agents make economic decisions, can be informative. In a previous study in 1988, the authors surveyed 2,030 households in four markets who had recently bought homes; three of those markets—Los Angeles, San Francisco, and Boston—were booming at that time, and the fourth, Milwaukee, was stable. For the present paper they essentially repeated the survey in the same markets in early 2003. The first three markets (which the authors call "glamour" markets) have experienced similar housing price cycles over the last twenty years. All three had a housing boom in the late 1980s or early 1990s, with prices increasing between 1982 and the first peak (1988 in Boston, 1989 in Los Angeles and San Francisco) by more than 125 percent. The bust that followed was most severe and longest lasting in Los Angeles, where prices fell 29 percent to a trough in early 1996, but price declines were also substantial in the other two. All three have seen a prolonged boom since their trough, with prices rising by 129 percent in San Francisco, 94 percent in Los Angeles, and 126 percent in Boston. Housing price behavior in Milwaukee could hardly be more different, with a steady climb of roughly 5½ percent a year over the period since 1982, in line with the growth of income per capita in that metropolitan area. Interestingly, over that entire period, home prices in Milwaukee, although starting and ending substantially lower than in the glamour cities, rose by roughly the same percentage as theirs did.

What caused the recent dramatic rise in prices in the glamour markets? The authors recognize that the substantial decline in mortgage rates may be part of the story. Because of the lower rates, despite the large increases in prices over the period, the carrying cost of a mortgage for 80 percent of market value is the same or lower today than in 1995 in all three markets. But they note that these mortgage rate reductions did not have a comparable effect on prices in Milwaukee, and that their regressions that included mortgage rates still greatly underpredicted the recent price increases in the glamour markets. In the 1988 survey the authors saw strong evidence of bubble psychology in the glamour markets: recent homebuyers said then that they were influenced by an investment motive, that they expected extraordinarily high future rates of home price inflation, and that they perceived little risk. They inferred that emotion and casual word of mouth played a significant role in home purchase decisions. The 2003 survey was sent to 2,000 persons in the same four metropolitan areas who had bought homes between March and August 2002. The response rate was 35 percent, about 10 percentage points lower than in 1988. Over 90 percent of respondents were buying a single-family home as a primary residence. The proportion who were first-time purchasers ranged from about one-third in Los Angeles to over half in Milwaukee.

The survey was designed to shed light on aspects of homebuyers' views and behavior that might be relevant to whether or not there was a bubble in their housing market. In particular, respondents were asked about their investment motivation, including expectations of future price rises and the risk of a price collapse, about the extent to which the housing market was a topic of conversation and source of excitement among their friends and associates, and about their adherence to popular theories about housing markets that the authors regard as simplistic or even fallacious. The authors hoped to see, among other things, whether the bust that had intervened since the previous survey had left a clear imprint on current psychology.

Case and Shiller suggest that a defining characteristic of a housing bubble is a widespread view on the part of buyers that housing is largely an investment. They find that only a small proportion of respondents in the glamour markets (7.5 to 10.6 percent) said they had bought strictly for investment purposes (3 percent or fewer bought to rent to others). In these markets, which had experienced a bust in the early 1990s, the media had given much attention to the possibility of a bubble in 2003. So it is not surprising that these proportions were down (if only slightly) from the 1988 survey. Nonetheless, for the vast majority (from 86 to 92 percent) of buyers in all regions, the investment motive was either a "major" consideration or "in part" a consideration in their purchase. Interestingly, the total was highest in Milwaukee, the "stable" market. In the glamour markets the proportion of buyers that perceived a "great deal of " or "some" risk increased relative to 1988, most dramatically in San Francisco, where the proportion seeing a great deal of risk increased from 4.2 percent to 14.8 percent. Although in all cities the proportion that perceived either "major" or "some" risk was over 55 percent, the authors conclude that the perceived risk of price decline remains small, and that homebuyers in general do not perceive themselves to be in a bubble.

Because unrealistic expectations of future price increases are a defining feature of bubbles, the authors asked a series of questions exploring these expectations. A high percentage of respondents expected an increase in home prices in the next several years. The median expected increases for the next twelve months were 10 percent in Los Angeles, 7 percent in San Francisco, and 5 percent in Boston and Milwaukee; the average expected increases were noticeably higher. However, expectations for the next ten years were far rosier. The average annual price increase expected in Los Angeles, San Francisco, and Boston was more than 13 percent, and in Milwaukee it was almost 12 percent. One wonders whether buyers understand the power of compounding: even in Milwaukee the reported expectations imply a tripling of home prices in ten years. Although the responses in the glamour cities for the next twelve months were lower than they had been in 1988, the ten-year expectations were even higher. Fewer respondents in 2003 said that now was a good time to buy a home because prices may be rising in the future, but at least two-thirds agreed with the statement in all four cities, and many thought there was a risk that delay might mean not being able to afford a home later. The authors also find that high percentages of respondents admitted to being influenced by "excitement" about home prices and to discussing the housing market with friends and associates. The authors judge that these indicators of bubbles were fairly strong in 2003, although generally weaker in the glamour cities than in 1988, and somewhat stronger in Milwaukee.

The authors explore respondents' agreement with a number of simple popular theories or stories about booming home prices for further evidence of a bubble mentality. They find widespread agreement with the statement that "housing prices have boomed in [city] because lots of people want to live here." This suggests that respondents confused a high level of prices with a high rate of change, but the authors recognize that respondents may have simply misconstrued the statement. More telling is that 20 to 40 percent agreed with the statement that "when there is simply not enough housing available, price becomes unimportant" and that, for more than half the respondents in the glamour cities, the best explanation of sellers getting more than one offer on the day they list a property, some of them over the asking price, was that "there is panic buying and price becomes irrelevant." Although these responses suggest that psychology is a major element in booming markets, the authors report that their respondents did not hold this view. Fifteen percent or fewer said they believed that the recent price increases reflected the psychology of homebuyers and sellers, and more than 80 percent thought that "economic or demographic conditions such as population changes, changes in interest rates or employment" were responsible. These results, together with the fact that respondents seldom mentioned psychological factors to explain recent home price changes, provide evidence that they did not believe themselves to be in a bubble.

Some observers have linked the beginning of the recent real estate boom to the stock market collapse that took place at about the same time. In theory, a stock market collapse could have two opposing effects: a wealth effect that reduces the demand for other assets, including housing, and a substitution effect, with a flight to assets of greater perceived quality, that increases demand for houses. The authors asked several questions to clarify the stock market's role. The vast majority of buyers stated that the stock market "had no effect on my decision to buy my house." Between a quarter and a third did state that the market decline "encouraged" them to buy a home, and only a very small percentage found it discouraging. Although the authors view the evidence as inconclusive, they note that many related comments by investors suggested that they had lost their appetite for stocks with the recent market decline and its volatility and had come to view real estate more favorably.

A substantial number of the buyers in the survey were also sellers, and the authors questioned them to learn more about the apparent upward rigidity of asking prices in boom markets. Forty-five percent of the respondents in the San Francisco area reported selling above the asking price in 2002. The shares in the other three cities were close to 20 percent. Many sellers thought they would have sold just as quickly if they had charged 5 to 10 percent more. When asked why they did not ask more, a surprising number of sellers—a majority in San Francisco and Boston, a near majority in Milwaukee, and 26 percent in Los Angeles—reported that it would be unfair, and many said that "the property simply wasn't worth that much." The authors also find sellers' attitudes to be consistent with the idea that asking prices are rigid downward.

The authors conclude with their own description of speculative bubbles in housing. They recognize that supply and demand for housing are affected by a large number of fundamental factors, such as demographics, income, employment, interest rates, and construction costs. Changes in one or more of these fundamentals can increase demand, putting upward pressure on prices. But they believe that expectations can amplify these effects when supply is inelastic. Price increases caused by changes in fundamentals can lead buyers to anticipate continued increases in prices, further increasing demand and setting off an upward spiral. In this view the long economic expansion of the 1980s drove up demand in the boom-bust cities. In the short run the increased demand encountered inelastic supply, inventories of properties for sale shrank, and vacancies declined, setting off the boom. As happened in that instance, longer-run forces eventually reverse the impact of the initial increase in demand and the public's overreaction to it. In some markets new construction can bring new housing on line in a relatively short time. In other markets, where supply is less elastic because of a fixed supply of land, or because of zoning or other restrictions, the price increases themselves are eventually self-limiting, in that they make living in the area less desirable, decreasing the labor supply and discouraging industry.

Is the present boom in housing prices a bubble, and is there a risk of a broad collapse that could adversely affect the national economy? The authors see little risk of such an outcome. They believe the fundamentals can explain the home price increases observed in most of the United States since 1995. For more than forty states, income growth alone explains virtually the entire increase in housing prices, and, in the vast majority of states, falling interest rates have actually made housing more affordable than it was in 1995. However, the authors do believe there is substantial evidence of a bubble psychology at work—high expectations of future price increases, together with a sense of opportunity, urgency, and excitement—in the three glamour cities they studied. Although these indicators are not as strong as they were in the 1980s, and most people in these areas do not perceive themselves to be in the middle of a bubble, the authors note that most did not perceive a bubble in 1988 before that bubble burst. The authors believe it is reasonable to suppose that in some of these glamour cities price increases will stall, and perhaps even decline. But only in the unlikely event that such price declines are synchronous and spread to markets in which there is no evidence of a bubble could they become a significant drag on the national economy.