Editors' Summary

THE BROOKINGS PANEL ON Economic Activity held its eighty-first conference in Washington, D.C., on March 30 and 31, 2006. This issue of the Brookings Papers includes the papers and discussions presented at the conference. The first paper takes a new approach to assessing the boom in home prices, using a model that parallels the one commonly used to value assets such as stocks. The second analyzes labor force participation and its determinants and projects future labor force growth. The third examines changes in wealth by age group and relates them to changes in law and the economy and to demographic characteristics. The fourth examines the present defined-benefits pension system and considers how to reform its regulation and insurance by the federal government.

IN FINANCIAL MARKETS THE hallmark of a bubble is an asset that is priced far above its fundamental value, which depends on the discounted stream of future cash flows—earnings, dividends, or interest. Differences of opinion about whether or not a bubble exists reflect differences of opinion about the fundamental value of the asset. In the case of owner-occupied housing, there is no readily available, easily estimated analogue to cash flow, and therefore opinions about fundamental value can differ widely. Many observers and market participants have instead focused on the behavior of prices themselves as a way of assessing whether "irrational exuberance" exists in the housing market. During the past five years housing prices have more than doubled in some metropolitan markets and have risen by 50 percent for the United States overall, leading many to conclude that there is a speculative bubble in housing. Those holding this view cite as evidence the sheer magnitude of the price increases, the rise in the ratio of average housing prices to average income, and the more rapid growth of home prices than of rents. In the first article of this volume, Margaret Hwang Smith and Gary Smith argue that all of these are flawed indicators of a housing bubble because they do not measure prices relative to fundamental values. Using data they have collected on individual homes in ten metropolitan areas, they calculate such a measure from the capitalized value of the stream of services that a home provides. They conclude that, in nearly all the markets they study, home prices remain near or below fundamental values.

The authors note that many of the features associated with a bubble as the term is commonly used—prices rising rapidly, a speculative focus on future price increases rather than the asset's cash flow, and the likelihood of an eventual price collapse—could also be present in a bubble defined relative to fundamental values. But not necessarily. Prices can increase rapidly for a considerable period and still remain below fundamental values, as they are likely to do when a significant, unexpected drop in mortgage rates both increases fundamental values and leads to rapid price increases. The authors stress that housing prices are not determined in a smoothly functioning, efficient market rooted in fundamentals. Instead they are influenced importantly by "comps," or the prices of recently purchased comparable homes; appraisers, real estate agents, and buyers and sellers themselves all rely heavily on data from these "comps." In this situation, speculative behavior can lead to rapid price increases or to price declines, and the corrective pull of fundamentals may be very weak. Hence the authors see their calculations as most directly relevant to individuals deciding whether to buy or rent, rather than to those deciding whether to buy or sell now rather than later. They themselves believe that short-run price movements are hard to predict, irrespective of the relationship between current prices and fundamental values, so that market timing is quite risky.

The authors provide an extensive discussion of what they call "bubblemetrics," arguing that most of the measures commonly used to label the current housing situation a bubble are flawed. These include the pervasive use of aggregate measures that ignore the heterogeneity of homes, differences in location and quality, and the use of overly simple measures of affordability. They note that the conclusions of more sophisticated regression models depend on the implicit assumption that the historical samples of home prices used in the analysis are randomly distributed around fundamental values. Given that assumption, systematic deviations of current prices from predictions generated by these regressions can be regarded as evidence that prices have wandered away from fundamental values. But the authors reject that assumption and note that past prices may have been too low, so that recent price increases may have simply brought them more in line with fundamentals. They are also critical of studies that simply look at the ratio of prices to rents or compare their rates of growth, because these ignore the many factors that should in principle affect the price-to-rent ratio, including interest rates, tax laws, and the nonlinear relationships among these variables. They also note that the dwellings sampled in rent indices are typically dissimilar to those sampled in sale price indices, so that the comparison becomes one of apples to oranges.

In the authors' model the basic measures used to evaluate whether housing is overpriced are net present value (NPV), fundamental value (P^*), and the internal rate of return (IRR). The NPV is the sum of the present values of all "cash flows," including the purchase price and associated mortgage costs, on the minus side, and the value of the housing services provided, on the plus side. When the NPV is positive, purchasing a home is preferable to renting, and when negative, renting is the better choice. The fundamental value (P^*) is simply the price that would imply a zero NPV. The premium is the excess of P^* over the actual price (P), expressed as a percentage. The required rate of return used in these calculations to discount future flows should in principle depend on rates of return, adjusted for risk, on the investments forgone because of the home purchase. The IRR is a hypothetical annual discount rate that would make the NPV and the premium zero.

The authors present simple models showing the algebra of present value calculations and the relationship among these variables. Calculations using these models make it clear that the relationships between the NPV, the discount rate, and the growth rate of net cash flow are highly nonlinear. These nonlinearities imply that many empirical models are misspecified. For example, for long horizons, with zero growth in cash flows, a reduction in the discount rate from 6 percent to 4 percent increases the present value by approximately half; yet with a growth rate in net cash flow of 2 percent, the same change in rates doubles the present value. The simple model also shows how a mortgage, by leveraging the homeowner's equity, increases the fundamental value and the rate of return so long as the mortgage rate is less than the homeowner's required, unleveraged rate of return. The realistic calculations that the authors later make for their sample are more complicated than those just described, because they account for substantial transaction costs, various elements of cash flow that do not necessarily grow at the same rate, and the fact that conventional mortgages are amortized, which implies that loan balances and tax-deductible interest costs decline over time.

How does one measure the net cash flow required to compute fundamental value when the housing services for owner-occupied housing are not priced? The authors do this by assuming that buying and renting comparable homes are often viable alternatives. Hence, for their empirical work, they use the observed cost of renting a home as a measure of the services received by the owner-occupier of a comparable home—this rental cost that the owner-occupier avoids is the parallel to gross cash flow produced by typical investment assets. The authors allow for future changes in rents by calculating them for a range of realistic rental growth rates. They acknowledge the existence (but do not take account) of other differences between renting and owning, such as control over furnishings and décor, greater benefits from investments in improvements, and greater privacy; all these factors suggest that using only avoided rent to measure the benefits of ownership is likely to understate fundamental value.

Converting gross to net cash flow requires subtracting a variety of costs. The authors make an extraordinary effort to use realistic estimates of all the important costs of home ownership—including brokerage fees, closing costs, maintenance, insurance, property taxes, and mortgage costs—and to allow for the housing subsidies implicit in state and federal income tax policy. Estimating mortgage costs requires assumptions about interest rates, points (a percentage of the sale price assessed by the lender at closing), and other terms of a mortgage. Since property taxes, mortgage interest, and points are tax-deductible, the authors also take account of differences in state and local tax rates.

For their empirical analysis, the authors assemble rent and sale price data on carefully matched pairs of rented and owner-occupied single-family homes. Using data for homes purchased in the summer of 2005, they match these homes with rental homes that are similar in terms of size, amenities, and location. In a relatively small number of cases they find exact matches, where an owner-occupied home was bought and then rented, or vice versa. Data were collected in ten metropolitan areas, including some where price increases had been among the highest in the nation, such as Boston, Los Angeles, and San Francisco, and others where the increases had been below average, such as Dallas, Indianapolis, and New Orleans. In some areas that are large and geographically varied, such as Los Angeles, they limit their study to relatively homogeneous neighborhoods.

For each purchased home in their sample, the authors compute the NPV and other measures for a variety of assumptions about discount rates, future growth in rents on the matched rental property, and how long the buyer of a home expects to own it. For the base case they assume a 3 percent annual increase in housing rents and expenses and a 6 percent required after-tax rate of return, a 20 percent down payment, a thirty-year mortgage with a 5.7 percent fixed annual mortgage rate (the average thirty-year rate in mid-July 2005), closing costs equal to 0.5 percent of the sale price, annual maintenance costs equal to 1 percent of the sale price, and a 6 percent transaction cost to the seller if the home is sold. The federal marginal income tax rate is assumed to be 28 percent, and the capital gains rate is assumed to be 15 percent on gains above $500,000. State and metropolitan-area data are used for property taxes, state income taxes, and home insurance.

Because of the significant transaction costs involved in buying and selling a home, the NPV and the IRR are sensitive to the length of the holding period, particularly for short horizons. Although the authors compute these variables for a wide range of holding periods, they focus on two cases: an infinite horizon and a ten-year horizon. They observe that the infinite horizon is more relevant than it might seem at first and avoids the need to assume a selling price. In particular, for horizons of twenty or thirty years, transaction costs are of minor importance, and for a homeowner who contemplates changing homes but remaining in the same metropolitan area, changes in local home prices may be of minor importance: selling high means buying high, and vice versa. In the base case for finite horizons, it is assumed that the price at the time of sale has grown at 3 percent a year.

The results of the calculations are in striking contrast to widespread opinion. Although, on average, prices on homes in San Mateo County, California (just south of San Francisco), greatly exceed fundamental values, in the other markets average prices are either roughly in line with or below fundamental values. For a homebuyer in San Mateo with a discount rate of 6 percent, the average home is 54 percent overvalued at an infinite horizon, and 42 percent overvalued at a ten-year horizon. The corresponding IRRs are well below the required rate of return. However, prices in Orange County, California, are roughly in line with fundamentals, and homes in Los Angeles, San Bernardino County, Boston, and Chicago appear somewhat underpriced. Home prices in Atlanta, Dallas, Indianapolis, and New Orleans are substantially below their calculated fundamental values. In Indianapolis, where the average rent is about half what it is in Boston and sale prices are about one-fourth, the authors' calculations show fundamental values to be roughly triple market prices. In other words, in all but one of the markets they investigate, the authors calculate that housing prices are near or below (sometimes well below) fundamental values; purchasing a home does not need to be rationalized by a belief in rapid and unsustainable further price increases.

The use of individual matched pairs enables the authors to estimate the distribution of price premiums, not just the mean. Not surprisingly, there is substantial variance in these premiums. Although it is hard to find a bargain in San Mateo, in Orange County there are almost as many homes selling below fundamental value as selling above.

Estimates of fundamental value depend crucially on assumed growth rates for rents and, for short horizons, growth rates of prices. The authors check the robustness of their base case results for ten-year and infinite horizons, calculating the values for annual growth rates in rent of 2, 3, and 4 percent (which cover the range of actual growth rates across the ten metropolitan areas during 1985–2005), and for growth rates in sale prices of 0, 3, and 6 percent. They show that the extremes imply large differences from the base case, but that for most variations the qualitative conclusions hold. If the annual growth rate of rents were permanently 4 percent—approximately the growth rate in the historical period in Boston, Chicago, and the four California markets—home prices in all these areas, even San Mateo, would be financially justified with an infinite horizon, and prices in all areas except San Mateo and Orange County would be justified with annual rent increases of 3 percent. Prices in Atlanta, Dallas, Indianapolis, and New Orleans are justified even with an annual growth rate of rents of only 2 percent. In Indianapolis the IRR exceeds 13 percent for both horizons even assuming no growth in either rents or prices.

Although uncertainty about growth in rents is the major risk for a homebuyer with a long horizon who takes out a thirty-year fixed-rate mortgage, uncertainty about the future sale price becomes more important as the horizon shortens. The authors provide schedules showing the IRR for the full spectrum of horizons and a variety of rates of price growth. These show, for example, that for a ten-year horizon, if prices do not increase at all and rents grow at only 2 percent a year, only four of the ten markets have prices below fundamental values, and homes are overpriced by more than 30 percent in all the others.

The authors further explore the risks from uncertainty about the growth rates of prices and rents by conducting Monte Carlo simulations for a representative house in the Los Angeles area, using stochastic processes for rents and prices estimated for the period 1983–2004. The risks are greatest for intermediate horizons and are quite substantial. The authors choose as their representative home one that has a sale price of $571,098 and an NPV of $25,539 under the base case assumptions. For a ten-year horizon the simulations indicate a 16 percent chance that the NPV will be –$50,000 or less, and a 4 percent chance that it will be –$100,000 or less. The probability of these extreme negative outcomes is roughly offset by that of extreme positive outcomes; median NPV is trivially different from that in the base case.

Changes in mortgage rates have very large effects on fundamental values. Homeowners with long-term, fixed-rate mortgages are effectively insulated from the risks of such changes until they sell, and indeed they benefit from the option to refinance if rates fall. However, uncertainty about rates increases the risks to homeowners with variable-rate mortgages. Monte Carlo simulations using an estimated stochastic process for rates show that this added risk is worth taking into account but is small relative to the risks associated with growth of rents and prices. Of course, the fundamental value of the home at the time of sale, and its sale price if prices move with fundamentals, are significantly affected by future rate changes.

The possibility that Congress might change the tax treatment of mortgage interest, as proposed, for example, by the President's Advisory Panel on Federal Tax Reform in 2005, poses another risk to homeowners. The authors calculate that, for a ten-year horizon, replacing current deductibility with a 15 percent tax credit, even without a cap on the credit of the type proposed, would reduce the median estimates of fundamental value across cities by 11 to 17 percent, and complete elimination of the tax preference would reduce it by 20 to 30 percent.

If the authors' central findings are correct, what explains the extraordinary increases in housing prices over the last five years, and how can they fail to be far above fundamental values after such a rise? In the authors' model, one way is that fundamentals may also have grown rapidly over the period as a result of some combination of mortgage and discount rate reductions and growth in rents. Another is that housing prices were low relative to fundamentals at the beginning of the period and have more or less caught up. Determining the relative importance of these two factors is difficult given the lack of historical information on prices and rents. The authors manage to gather matched rental and sale data in the Los Angeles metropolitan area for 2001–04. Calculating fundamental values as before, but using historical values for the model's key parameters, they find that home prices three to five years ago were roughly 40 percent below fundamental values. In 2004 this gap between fundamental values and prices was more than halved, falling to a level roughly in line with their results for 2005.

To extend their calculations to earlier years, for which they do not have data on individual homes, the authors use the housing price index constructed by the Office of Federal Housing Enterprise Oversight and the Bureau of Labor Statistics' owner's equivalent rent index to backcast to 1983 the annual prices and rentals for a home with the average characteristics of their 2001 sample. Their results are again striking. In 1983, even with a thirty-year mortgage rate of roughly 13 percent and what was then a conservative 3 percent expected growth rate of rents, sale prices were 14 percent below fundamental values. In 1984, with the rise in mortgage rates to 14.7 percent, this discount fell to 4 percent. Then, as mortgage rates started falling, the discount at first rose rapidly and then disappeared with the near doubling of home prices between 1985 and 1990. With the continuing decline in rates and the decline in home prices in the early 1990s, prices again fell below fundamental values. Despite rapid growth in sale prices after 1996, the gap remained large, shrinking to less than 10 percent only after 2003, when mortgage rates stopped falling. The authors observe that for the entire eleven-year period 1993–2003, buying clearly dominated renting for households with long horizons.

Forecasts of the change in the price premium from 2001 to 2005, derived using the price and rent indices from the above exercise, closely follow the authors' estimated change based on their data from matched pairs. This suggests the usefulness of using the same method for backcasting from the 2005 matched-pair data for their other markets. The pattern of intertemporal changes in fundamental value is quite similar across markets, largely reflecting the fact that changes in interest rates and federal taxes are common to all. However, the changes in market prices are quite different: San Mateo prices are quite volatile and more often above than below fundamental values. The patterns for San Bernardino County, Orange County, and Los Angeles are quite similar, with prices generally below fundamental values at first, then a run-up in the late 1980s, followed by a falling off and then an eventual closing of the gap. In Boston, Chicago, and Dallas, sale prices were close to fundamental values for much of the 1980s and have lagged behind since. In Atlanta and Indianapolis, prices remained well below fundamental values over the entire twenty-three-year period.

The authors conclude by observing that if the essential feature of a bubble is that prices rise far above fundamental values, there was no bubble in the prices of single-family homes in 2005. However, they do not claim that housing prices are always aligned with fundamental values, and indeed they suggest it would take a very peculiar set of assumptions to place fundamental values in the middle of the fluctuating market prices of the past twenty-three years in any of the ten markets they examine. They believe this helps explain how, despite rapidly rising prices, prices can be still below or near fundamental values in many urban areas. For potential buyers who plan to stay in the same area for many years, the relevant question is not whether prices have risen rapidly, nor even how fast they expect them to rise in the future, but whether at current prices the home is a fundamentally sound investment. For the typical home in nine of the authors' ten markets, their answer is yes.

THE LABOR FORCE PARTICIPATION rate measures the percentage of the working-age population that is in the labor force, either employed or looking for work. Understanding variations in participation rates is useful both for analyzing the economy's long-run growth prospects and for interpreting the degree of slack or tightness in labor markets at a point in time. Over a business cycle, participation rises when jobs are plentiful and falls when jobs are hard to find. Over longer intervals, the trend in participation is affected by changes in the age distribution of the population and by policy and societal changes that affect people's desire and incentives to work and the hiring practices of firms. The participation rate rose almost without interruption from the mid-1960s to the late 1980s, mainly as a result of the large number of women joining the labor force. After falling in the recession of the early 1990s and recovering in the subsequent expansion, participation fell noticeably between 2000 and 2005. In the second article in this volume, Stephanie Aaronson, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William Wascher examine the variations in the participation rate over the past several decades in order to identify what has moved it in the past, how to interpret its recent decline, and what to expect going forward.

Labor force participation varies predictably over the life cycle. The working-age population is, by convention, defined as persons 16 years and older, and participation rates are lower for the youngest and the oldest age groups than for those in between, with ages 25–54 being the peak participation years. This general pattern is observed for both men and women, and the typical variation over the life cycle is large enough that the demographic bubble of the baby-boom cohorts has had a pronounced effect on overall participation rates as those cohorts have moved through the life cycle. While conforming to this general pattern, actual participation for each age group has trended higher or lower for extended periods in response to economic and societal changes that impact some age-sex groups differently than others. The authors capture these features in a regression model that allows them to decompose the participation of each age-sex group into four parts: a typical age profile of participation; an average birth cohort effect that shifts this age profile according to when individuals were born; the separate effects of a number of "structural" variables that capture economic and social changes that influence different age groups (or, equivalently, different birth cohorts) at various points in time; and business cycle effects.

The authors assemble quarterly data starting in 1977 that are disaggregated by age and sex. The data come from the monthly Current Population Survey along with its Annual Demographic Supplement and include unpublished micro data. The authors divide the observations into fourteen age categories, ranging from 16- and 17-year-olds to persons aged 70 and older; they model males and females separately, resulting in twenty-eight separate regression equations, each explaining participation rates for a single age-sex group. The authors sketch the evolution of various economic and social changes that are likely to have had an effect on participation, and they select time-varying structural variables to capture such effects in their regressions. They organize their variables into three categories: human capital, financing of nonparticipation, and family structure. Human capital variables include, for ages 16–24, returns to a college education; for older age groups they include average educational attainment when a cohort was 27 years old; and for 18- to 61-year-old women they include the female-male gap in weekly earnings. Variables measuring financing of nonparticipation include sex-specific life expectancy at age 65 (as a measure of the length of retirement to be financed) and various changes to Social Security's old-age and disability insurance provisions. (Some measures of wealth were also tried but were found to have no explanatory power.) Family structure variables are the presence of children under age 6 and the percentage of women (in each age group between 18 and 61) who are married. Business cycle effects are captured by deviations of output from a trend generated using a Hodrick-Prescott filter.

The regressions find significant effects of the business cycle on labor force participation for most age-sex groups. The effects are positive and large—lowering participation in slack economies and raising it in booms—for the youngest (ages 16–24) and the oldest males (ages 60 and over) and for females up to age 44. The effects are small for 35- to 59-year-old men, indicating very strong attachment to the labor force in those years, and, interestingly, they are negative for 55- to 64-year-old women. For their structural variables, the authors report the effect on aggregate participation of a 1-standard-deviation change in each. Since most are relevant only to particular age-sex groups, most of the effects calculated in this way are modest. Aggregate participation is moved by 0.1 percentage point or more by a 1-standard-deviation change in returns to education, eventual college attainment, life expectancy, the female-male earnings ratio, the marriage rate, and the presence of children under 6. The largest effect is a negative 1.5 percentage points from women being married, and the second largest is 0.4 percentage point from having graduated from college. As the authors observe, the inevitable imprecision of their variables in capturing the social and economic changes that matter, including the difficulty of properly aligning their timing, affects these estimates. Nonetheless, along with the cyclical variable and the average cohort and age effects, the regressions track the actual movements of participation reasonably well, with only the youngest and the oldest age groups showing regression errors as large as 1 percentage point.

By eliminating cyclical effects, the authors produce estimated participation trends over time for each age group. (For the youngest there are not enough observations to allow separation of cohort and cyclical effects, and so cohort effects are extrapolated from earlier cohorts.) For most male age groups, participation trends have declined fairly steadily over the past thirty years; the exception is older men, whose participation trend has risen since the 1980s. Trends among females have been far more volatile. The over-55 groups have had rising trends since the 1980s, those aged 25–54 had rising trends until the past ten or fifteen years and have seen some decline since, and the youngest groups have had a generally declining trend since the early 1980s.

Aggregating these results by weighting individual trends by population weights in each year, the authors estimate trends in participation through 2005 for the total population and for males and females separately. These broader trends are importantly influenced by the typical life-cycle pattern of participation and the shifting age distribution of the population over time. The estimated aggregate trend peaked in 2002 and has declined noticeably since. More than half of this decline is due to the aging of the baby-boomers, who are beginning to move from high- to low-participation age groups. Although there is uncertainty about other determinants of future participation, such as the number and age composition of immigrants and changes in participation within age groups, this ongoing demographic shift toward an aging population will almost surely dominate, producing declining participation rates for decades.

Going beyond this general characterization of the future, the authors use their disaggregated analysis to provide a quantitative range of projections for participation rates. These make use of the separate identification of age (life-cycle) effects, fixed birth cohort effects, and time-varying effects from structural variables in their model. In their baseline projection, for cohorts already 16 and older at the end of their sample period, they hold cohort effects constant and age these cohorts according to their latest estimates of age (life-cycle) effects. For their structural variables they allow life expectancy to evolve in line with Census Bureau projections and assume that educational attainment for cohorts currently younger than 27 changes at the same rate as that of recent cohorts. Other structural variables are held at their end-of-sample values. For cohorts entering the labor force after 2005, the authors keep fixed cohort effects at the average value of recent cohorts. Using these assumptions and official Census Bureau projections of the population including immigrants, the authors project a steady downward trend in participation that reduces the aggregate participation rate by more than 3 percentage points by 2015. By comparison, this is approximately how much participation rose during the 1970s, the decade of the steepest increase in women's participation. About 2 percentage points of the decline comes from the aging of the baby-boomers, with the rest coming from the model's prediction of the change in participation within age-sex groups.

The authors consider some plausible departures from this baseline projection. Model errors at the end of their sample period are small for aggregate participation. However, in the last few years, participation was well below predicted values among 16- to 19-year-olds and well above predicted values for workers aged 62 and older. Treating these errors as age group related has a negligible effect on overall participation projections for 2015, with the higher participation for older workers adding about ¼ percentage point, and the lower participation for the youngest workers subtracting about the same amount. If instead the recent model errors for the youngest workers are treated as a cohort effect for this and subsequent cohorts, the effect on the 2015 projection is much greater, reducing participation by over 1 percentage point in that year and by growing amounts thereafter. The authors also point out that projections of labor force participation from their model would be sensitive to changes in their structural variables, and they consider the effect that improved educational attainment would have. They note that school enrollment rates for teenagers have risen from about 60 percent in the late 1980s to nearly 75 percent in 2005. Although these higher enrollment rates help explain the declining participation of the youngest workers, the authors find it plausible that they may lead to higher rates of college attainment as these cohorts mature. Estimating this effect by holding cohort effects constant at the 1984 level for cohorts born after 1984 (rather than letting them decline as in their baseline projection) raises 2015 aggregate participation by ½ percentage point, and by more thereafter.

The authors go beyond their model to consider further evidence on the nature of changes in participation in the most recent years, the period when it is most difficult to distinguish transitory cyclical variations from changes in trends. This period is also the most important for projections, since future trends take off from the estimated trend at the end of the sample. To examine this post-2000 period more closely, the authors conduct a cross-state analysis of participation, using data spanning 1990–2005 and allowing for a break in the trend and in cyclical sensitivity starting in 2001. Their statistical analysis relies on the nonsynchronous fluctuations of labor demand across states to identify the portion of participation changes that were unrelated to such fluctuations and therefore represented trend rather than cyclical effects. This analysis shows that cyclical and structural factors each accounted for about half the post-2000 decline in participation, with little change in cyclical sensitivity. This result is quite similar to the estimates from the cohort-based time-series model reported above. The authors obtain further insight on this period using data on gross labor force flows to compare the most recent recession and recovery with predictions based on historical estimates of the cyclical response. This exercise shows that flows from employment to nonparticipation were much larger than predicted in recent years, a result that is inconsistent with a weak labor market, in which workers would be unusually worried about finding a new job. This finding further supports the importance of structural rather than cyclical forces in the decline in participation in this period.

The economy's total labor input depends not only on participation rates but also on the size of the working-age population and on average hours worked. The authors turn briefly to these other two factors to round out their projections of the future trend in labor input. Over the past decade, population growth has averaged 1.2 percent a year, and the Census Bureau projects a decline in this growth rate to about 0.9 percent a year by 2010 and 0.8 percent a year by 2015. The authors note that uncertainty about net immigration is important in these projections and that estimates by the Congressional Budget Office (CBO) and the Social Security Administration (SSA) assume higher immigration flows than the Census projects. Average hours worked, as measured in the widely used Current Employment Statistics, trended down noticeably until the early 1980s and more gradually thereafter. The authors discuss the several factors known to have influenced the trend in hours, including changes in industrial composition and demographic changes. Using a Kalman filter model with controls for cyclical conditions to calculate a trend for average hours, they estimate a trend decline between 2000 and 2005 of 0.2 percent a year.

In their central projection, the authors estimate an aggregate participation rate in 2015 of 62.5 percent, compared with projections of 65.0 percent from the CBO and 65.2 percent from the SSA. Combining their participation rates with Census population projections leads them to project a substantial slowdown in labor force growth over this period. The CBO and the SSA also project slowdowns, but not as severe. In 2015 the authors project the labor force will be only 4.0 percent larger than in 2005. This is 4.4 percentage points below the SSA's projections and 3.7 percentage points below the CBO's. Using the authors' projections, if average hours worked were to continue to decline at the 0.2 percent annual trend rate they estimate for 2000 to 2005, total labor input would not be growing at all by 2015.

The authors caution that all these estimates are subject to considerable uncertainty. They note the difficulty, discussed above, of capturing the reasons for the recent participation behavior of the youngest and the oldest workers and the uncertainty that this creates for projections. They also observe that not only is it difficult to project the number of immigrants who will enter the country over the next decade, but the number that materializes could affect the aggregate participation rate, since immigrants tend to have higher than average participation rates. However, while emphasizing the uncertainty around any point estimates, the authors conclude that demographic factors will almost certainly lead to a substantial slowing in labor force growth over the coming decade.

DURING THE DECADE OF the 1990s, the aggregate net worth of Americans doubled, from \$20 trillion to \$42 trillion, despite a low and falling national saving rate. Much of the increase in wealth reflected extraordinary gains in the stock market, whose value grew by \$12 trillion, and in housing during the latter part of the period. These events, along with more widespread diffusion of stock ownership and a substantial increase in participation in and contributions to defined-contribution pension plans, make this period an especially rich one in which to examine household saving and wealth accumulation. Previous research has examined the active saving of different cohorts during this period, the asset composition of the growth in wealth, and the extent to which households chose to use their additional accumulated capital gains to finance increased consumption or early retirement. These studies typically focused on tracking particular birth cohorts through time. In the third article of this volume, William Gale and Karen Pence take a different perspective, asking to what extent successive cohorts of American households are wealthier than their predecessors.

Focusing on the striking aggregate wealth increases of this period, Gale and Pence explore how the wealth of individuals of a given age in 2001 compares with that of individuals of the same age in 1989. Data from the 1989–2001 Surveys of Consumer Finances (SCF) enable them to examine not only differences in the level and composition of wealth for successive cohorts, but also to hold constant demographic characteristics such as marriage, health status, educational attainment, and time in the labor force. They conclude that the median wealth of individuals of a given age and set of demographic characteristics changed little over this period despite the large aggregate capital gains that occurred. Remarkably, for most age groups, differences in the demographic characteristics of successive cohorts appear to explain almost all of the difference in their median wealth.

The authors begin by showing the differences in total real wealth between 1989 and 2001 of successive cohorts at different ages, without taking into account differences in their demographics. They define six age groups (determined by the age of the head of household), with the youngest aged 25–34 and the oldest 75–84, and they obtain data for about 500 households in each group. They find no significant increase in median wealth between 1989 and 2001 for the first three age groups; indeed, the median net worth of households with a head between the ages of 35 and 44 actually fell from $108,000 in 1989 to $99,000 in 2001. Mean wealth does increase for all three of these groups, but by only modest amounts. In contrast, both median and mean wealth for the three older age groups increased substantially. For example, median wealth for households with a head between the ages of 65 and 74 increased by almost 60 percent over these twelve years. This pattern is not restricted to the middle of the wealth distribution but rather emerges across most of the distribution. However, the authors note also that the pattern of increases is not uniform: for example, for the 65- to 74-year-olds there are significant increases at the 30th and the 80th percentiles, but not the 90th. Plots of the cumulative distributions of wealth for the 65–74 and 35–44 age groups in the two years also indicate that not only the rich got richer: the distribution for the older age group shifts upward over the entire range, and the 2001 distribution for 35- to 44-year-olds lies roughly on top of that for 1989.

Are the increases in median wealth of the older age groups between the two years concentrated in particular asset categories—for example, financial assets—where the extraordinary stock market gains might be expected to show up? The SCF gives detail on the asset composition of wealth, enabling the authors to examine differences in retirement assets, other financial assets, home equity, and "other" assets (which includes equity in vehicles, investment real estate, and closely held businesses). The authors also extend the SCF definition of retirement assets to include all defined-contribution balances rather than only those under the household's control before retirement, and they make their own estimate of a household's claim on defined retirement benefits. They reason that excluding these assets underestimates wealth and biases estimates of the change in wealth upward because of the well-documented shift from defined-benefit plans to a variety of types of defined-contribution plans during the period.

In general, the changes in the age profile of average holdings of the various asset classes between 1989 and 2001 are similar to the changes in the age profile of total net worth, but somewhat noisier. For the three younger age groups there are only small differences in the holdings of most asset classes across the twelve-year period, whereas for the three older age groups the differences are substantial. Differences in the age profiles for retirement assets mimic the differences in average wealth. But the age profiles for the other assets have distinctive patterns, as do their differences between 1989 and 2001. In 1989 home equity was almost the same for the 45–54, 55–64, and 65–74 age groups. In 2001, by contrast, a substantial increase with age appears: home equity for the 65–74 age group was roughly 40 percent greater than that for the 45–54 age group. In 1989 financial assets rose almost linearly with age; in 2001 financial wealth behaved similarly for the two youngest age groups, increased much more rapidly for those in their middle years (more than doubling the 1989 level for the 55–64 age group), and leveled off for the two oldest groups. The age profiles for "other" wealth follow a similar pattern, but with the 75–84 age group's wealth almost identical in 1989 and 2001.

What explains the substantial gains in wealth of older households in 2001 compared with those of similar ages in 1989? An individual of a given age in 2001 is a member of a cohort born twelve years later than an individual of that age in 1989 and will have faced a different economic environment than the earlier cohort at each stage of life. Consider, for example, an individual who enters the labor force during a period of tight labor markets and high wages, or one who reaches middle age with substantial accumulated assets just as the stock market enters a boom. These individuals would be expected to have greater wealth, in middle age and later, than someone who began his or her career in a weak labor market, or who faced a declining stock market after accumulating a solid financial nest egg. But demographics matter for gains in wealth also. And even though the demographic characteristics of the population change slowly, substantial differences can occur over twelve years. Although the authors' ability to investigate the importance of differences in economic environment over the lifetime of a typical individual or household is limited by the short time period covered by their data, they do have detailed information about the demographic characteristics of each age group in the two years they analyze. This enables them to investigate empirically the extent to which changes in demographic characteristics alone explain the differences in wealth for individuals of a given age in those two years.

The authors utilize information on sex, marital status, health, education, and labor force participation. They note that each of these demographic characteristics might be expected to have significant effects on wealth and saving. Married households benefit from economies of scale; widowed households typically face reduced pension and Social Security income; healthy individuals may work more hours, spend more years in the labor force, and have lower medical expenses; the better educated tend to have better-paying jobs and to have better health, even controlling for income and wealth; and longer participation in the labor force is likely to increase lifetime earnings.

The authors find that several of these demographic characteristics "improved" for older households in 2001 compared with those in 1989, which is consistent with their greater wealth, and did not improve or actually deteriorated for younger households over the same period. For example, both the share of married household heads and the share with good health rose among older households, but among younger households the married share decreased (perhaps because of a greater tendency to delay marriage), and health, surprisingly, deteriorated. Perhaps because of relative improvement in the life expectancy of men, the share of households headed by a widow in the 65–74 age group fell almost by half during the period. Educational attainment rose substantially for females in all age groups, but it rose more for successive older male cohorts than for younger cohorts. Most notably, the share of men in the 65–74 age group with post-secondary education increased from 31 percent in 1989 to 49 percent in 2001, whereas the share fell from 60 percent to 56 percent in the 35–44 age group. Lifetime participation in the labor force increased for females of all age groups.

The authors first examine the quantitative importance of these changes between 1989 and 2001 in the demographic characteristics of the different cohorts by estimating wealth equations for each of the six age groups. They allow a different intercept for the two years and run the regressions with and without the demographic variables, which are constrained to have the same effect in each year when they are included. Thirteen demographic variables are considered: dummy variables are used to distinguish postsecondary education from less education, and to distinguish “fair or poor” from “excellent or good” health, entered separately for males and females; three dummies are used for marital status—married, second marriage, and divorced or separated (single or widowed is the omitted category); an additional dummy is included for female-headed households to capture gender-based differences in the impact of divorce or death of a spouse. Continuous variables are used for years married, years since divorce or death of a spouse, age of household head, and years working full time (entered separately for males and females). The authors estimate “median” regressions, minimizing the sum of absolute deviations for two specifications of functional form: one in absolute levels, and the other using a wealth transformation that is much like the logarithmic but avoids the problem of taking logarithms when some observations are negative or near zero. The results are qualitatively similar for the two specifications.

The regressions without the demographic variables support the conclusions drawn from the earlier plots of age profiles of wealth for the years 1989 and 2001. Median wealth for each of the three younger age groups is statistically indistinguishable across the two years. In contrast, the differences in median wealth are large and highly significant for the three older age groups: wealth is greater in 2001 than in 1989 by amounts ranging from \$66,000 for the 55–64 age group to \$95,000 for the 74–85 age group, in the regression using the absolute wealth specification. Inclusion of the demographic variables does not change the conclusion that there is no statistically significant difference for the younger age profiles between the two years. For the older ages the statistically significant differences remain when only certain subsets of the demographic variables are included, but not when all the demographic variables remain in the equation. It appears that changes in birth cohorts’ demographic characteristics explain virtually all the differences in median wealth between 1989 and 2001.

What are the estimated partial effects on wealth of the changes in particular demographic variables? The authors report coefficient estimates from median regressions pooling 1989 and 2001 data for two age groups, 35–44 and 65–74. For both age groups nine of the thirteen variables are statistically significant. Most of the coefficients are of meaningful size and have the expected sign. For both males and females, postsecondary education is “worth” over \$70,000 in added wealth for the younger age group and more than triple that for the older group. “Fair or poor” health reduces wealth, in the case of the 65–74 age group by \$79,000 and \$97,000 for men and women, respectively. Married households have substantially more wealth than single or divorced households. The effects of years working, years married, and years since divorce or death of a spouse are small, as are differences in age within the age group.

In the regressions discussed thus far, the demographic variables are constrained to have the same effect on wealth in 1989 and in 2001. In principle, changes in average wealth of a given age group across the period could be apportioned between differences in the coefficients (betas) on the demographic characteristics (that is, in the magnitude of the effect of a given characteristic on wealth) and changes in the characteristics themselves. To examine this possibility, the authors estimate separate ordinary least squares equations for the two years. Although the estimated coefficients differ, only a few of the differences are statistically significant. Algebraically, the change in mean wealth for a given age group can be decomposed into the differences in the betas for the two years times the average values of the demographic characteristics in one year, plus the differences in the average values of the demographic characteristics times the values of the betas in the other year. In conducting this decomposition, the authors focus only on the older age groups where there is a significant difference in average wealth for the two years. They find that about half of the wealth increase for the 55–64 age group and essentially all of the increase for the two oldest age groups can be attributed to changes in demographic characteristics. Decompositions based on the estimates using the nonlinear transformation of wealth described above show that, for all three older age groups, essentially all of the increase can be attributed to demographics.

Even though differences in demographics appear to explain the differences in the median and average wealth of individuals of the same age in different years, much of the variation in wealth among individuals with the same characteristics is inevitably unexplained in these regressions. Is the distribution of household wealth in 2001 similar to the 1989 distribution once changes in demographics are taken into account? The authors use two complicated procedures to address this question. The first effectively reweights the households in the 1989 sample so that they reflect the distribution of demographic characteristics in 2001. The actual wealth of a 1989 household with particular demographic characteristics is used rather than its predicted wealth from a median regression. In effect, the 2001 distribution is the weighted sum of the actual wealth in the 1989 sample of households sorted by their demographics, with weights that reflect the relative frequency of those demographics in 2001 rather than in 1989.

The second procedure estimates individual "centile" regressions on the 1989 data, with each regression placing different weights on the individual observations so as to best fit different centiles of the error distribution. The high-centile regressions, for example, place heavy weight on the extremely high wealth observations. As a result, the coefficients on the demographic variables are likely to be quite different from those estimated by a median or a mean regression. Although the high- and low-centile betas do a better job of "explaining" the extreme observations, and hence the entire distribution of wealth, it is difficult to give them a structural interpretation. In effect, variation in the centile betas substitutes for the large residual errors that remain using a median regression. The authors use a random draw from the estimated distribution of 1989 centile betas to forecast the wealth of a randomly selected 2001 household with a particular set of demographics. This procedure is repeated over and over again to generate the hypothetical wealth distribution for 2001. Because of the random selection of centile betas, 10 percent of the households with given demographics will have wealth forecasted using the top 10 percent centile betas, another 10 percent with those same demographics will have wealth forecasted using the 80 to 90 percent centile betas, and so forth.

The authors show that, with the exception of the 75–84 age group, both procedures generate 2001 wealth distributions very similar to those actually observed. The authors regard these results as the most persuasive evidence that demographic characteristics are a significant determinant of the greater wealth of the older households in 2001. However, it is difficult to know whether the similarity in the hypothetical and the actual 2001 distributions reflects anything more than similarity of the error distributions for the two years and the shift in the median implied by the median regressions—in other words, whether changes in demographics are responsible for the realistic appearance of the hypothetical 2001 wealth distribution.

All of the analysis thus far has focused on explaining the wealth of different birth cohorts reaching the same age in different calendar years. Tracking instead the experience of given birth cohorts over time might provide additional insight into the importance of capital gains in explaining increases in wealth over time. The SCF does not track a fixed panel of households. Therefore the authors examine the experience of a "synthetic" panel, comparing, for example, households aged 45–54 in 2001 with households who were between 33 and 42 years of age in the 1989 SCF. They start with the stock and residential real estate holdings of the cohort in 1989 and add an estimate of capital gains in each asset that accrued over the 1989–2001 period. They also allow for differences in mortality rates by education and by sex. They do not take into account any active saving during the period, but they remove bequests and exclude the value of privately held business, investment real estate, and defined-benefit wealth. Using these data, the authors find that the share of the change in wealth explained by capital gains increases as one moves through the successive 2001 age groups: It is about a fifth for the cohort aged 45–54 in 2001, a third for the cohort aged 55–64, and three-fifths for the cohort aged 65–74. For the oldest group the change in wealth is more than explained—implying that this group actively dissaved or received bequests, or both. The authors conclude that capital gains do not come close to explaining all the wealth accumulation for the cohorts likely to be in the accumulation stage of the life cycle.

The authors see their main results as serving "to highlight a long, but sometimes downgraded, tradition in economics—dating back at least to the original formulation of the life-cycle model by Franco Modigliani and coauthors—that emphasizes the role of demographic variables in wealth accumulation." They note that the same demographic factors that appear to have fueled the increase in wealth also are likely to raise the expenditure needed to maintain living standards in retirement. Married households consume more than single households, more highly educated households likely have had (and will want to sustain) higher consumption than less educated ones, and healthy households may have lower medical expenses but will also live longer, on average, and thus have to finance a longer retirement. The authors observe that the prospects may be particularly unfavorable for the baby-boom generation, who, despite having lived through the bull market of the 1990s, had no more wealth, on average, in 2001 than their 1989 counterparts. Based on their findings, the authors suggest that policies that raise investment in health, education, and other forms of human capital could have far-reaching consequences for saving and wealth accumulation.

DEFINED-BENEFIT PENSIONS ARE an important source of retirement security for many workers. Unlike defined-contribution plans, in which retirement benefits depend on the returns on the assets in the employee's own account, the benefits of defined-benefit (DB) plans are specified by contract and are the responsibility of the plan sponsor, generally the employing firm or, in a few cases, a group of firms. Although the relative importance of DB plans has been shrinking steadily, they still insure 44 million participants and, as of 2001, held $1.8 trillion in assets. Over thirty years ago, the Employee Retirement Income Security Act (ERISA) established rules governing DB plans and created the Public Benefit Guaranty Corporation (PBGC), funded by premiums paid by DB sponsors, to insure workers' promised benefits. By most measures this system has worked well. However, in recent years the financial position of the PBGC has deteriorated, as several large firms with underfunded plans have gone bankrupt and others appear likely to follow, posing a substantial risk of much higher claims on the PBGC in the future. These developments have raised questions about the future of the public insurance system and of DB plans themselves. In the final article of this volume, David Wilcox examines the rules governing DB plans and the PBGC, discusses the sources of the current problems and future risks, and proposes a comprehensive reform of the system.

As background to his analysis, Wilcox describes the principles underlying simple DB plans. Such plans promise covered workers an annuity determined by their employment history, and plan sponsors are required to support this promise through payments into the pension fund and the investment income from the fund's assets. If the market value of those assets matches the present value of the plan's future liabilities—and an ideal set of regulations would ensure that this is the case—the DB plan is fully funded.

In contrast to this ideal, Wilcox identifies a number of deficiencies in the present system and associates them with different types of risks to which they give rise. Current regulations allow firms to systematically underfund their obligations. When a plan sponsor fails, limits on the PBGC's guarantees may leave some workers with lower benefits than they were promised. PBGC insurance is mispriced in a way that encourages both firms and workers to take excessive risk. This mispricing leaves taxpayers and healthy firms to bear the risk of having to shore up the PBGC should it need additional funding to bail out failed plans. Investors have difficulty properly evaluating the financial condition of plans under current accounting rules, resulting in inefficient valuation of firms in financial markets. In addition, the regulations governing the present system are highly complex, resulting in added costs to firms with DB plans and compromising the ability of the PBGC to oversee them properly.

Wilcox discusses how ERISA rules often permit underfunding to arise and, having arisen, to be corrected only very gradually. He provides details of the complex regulations under which DB plans operate and the departures from an ideal system to which they give rise. Two particularly important examples illustrate how present rules contribute to the underfunding problem. In so-called flat-benefit plans (a common type of DB plan in which contractually promised benefits are not tied to wage increases but grow only in proportion with years of service), the only way of adjusting for nominal wage growth is through frequent revisions in the DB plan specifications. But under ERISA, any top-up funding required because of such revisions may be amortized over thirty years. As a result, current assets do not cover the present value of future liabilities even for some funds that are in legal compliance. One study found that flat-benefit plans accounted for 90 percent of total liabilities taken over by the PBGC. Shutdown benefits pose another significant problem. These arise from plant closings or permanent layoffs and may call for full benefits to be paid before the affected workers reach traditional retirement age. Yet the possibility that shutdown benefits will be paid is typically ignored in determining a plan's funding requirements, and the PBGC recently estimated that they could add $15 billion to its liabilities in coming years. Funding shortfalls also arise from fluctuations in the value of pension fund assets, particularly equities. Such shortfalls may not be reflected in fund accounting until several quarters after they occur, and sponsors are allowed an extended time beyond that to make up for them. As a striking example, Wilcox reports that, in September 2000, Bethlehem Steel had 73 percent of its DB plan assets invested in equities. The value of plan assets declined by 25 percent over the next year and by another 23 percent the year after that, at which time the plan was terminated.

Although the extent of underfunding provides a warning of potential claims on the PBGC, its liabilities formally materialize only when an underfunded plan is terminated. Wilcox notes that a sponsor may terminate its DB plan if it is in bankruptcy proceedings or has persuaded the PBGC that it must terminate in order to stay in business, and the PBGC itself may terminate a plan if it determines that its own liabilities will increase unreasonably if it does not. Wilcox also describes the circumstances under which the insurance obligation of the PBGC may not cover the full amount promised to all workers. For plans terminating in 2006, payments are limited to $47,659 for a participant beginning to draw benefits from the PBGC at age sixty-five. And the PBGC is required only to meet a fraction of benefit improvements implemented in the five years before a plan is terminated.

Wilcox also notes that the insurance premiums that the PBGC charges are not related either to the risk of the sponsor's business or to the risk of the assets in the fund's portfolio. The annual cost to sponsors is presently $30 per participant plus a variable-rate premium of $9 per $1,000 of unfunded liability. However, most unfunded liability is exempt from this variable premium for various reasons, so that even the observable risk in the funding balance is not reflected in premiums. These rates appear to be extremely low in light of historical experience with current regulations.

Before turning to the specific pension reforms he advocates, Wilcox presents three axioms that, he argues, should guide reform. The first axiom is that workers should bear no risk for the benefits promised to them. Since workers already bear own-firm risk, a riskless pension promise permits workers and their sponsor firms to reach a more efficient compensation arrangement. The second axiom is that taxpayers should not bear the burden of the risk involved in insuring pensions. The third is that low-risk sponsors should not have to cross-subsidize the insurance of high-risk sponsors. Such subsidies would tend to drive low-risk firms out of the DB plan system, leaving a riskier pool of plans for the PBGC to insure.

Wilcox divides the present problem confronting the PBGC into two parts: dealing with the overhang of existing and potential liabilities, and designing a system for the future that would satisfy his three axioms. Of the three options for dealing with the present overhang—taxpayers footing the bill, surviving sponsors paying economically unfair premiums, or the PBGC defaulting on its obligations—he favors the first on the grounds that a taxpayer bailout would make good on the government's promise as it is widely understood, without heightening the risk of an exodus from the DB system. For the future he notes that a new system would have to make changes of three sorts: tightening funding and portfolio rules, rationalizing the pricing of the PBGC's insurance, and improving the information provided about DB funds in firm's financial reports, in the government budget, and to workers.

Turning to his proposals for a new system, Wilcox first considers how plan sponsors' required contributions should be determined. Sponsors should be required to fund benefit accruals each year and to amortize any shortfalls much faster than at present, all on a mark-to-market basis. Meeting this requirement would require the following main changes, which would also simplify current law: First, a single measure of benefit accruals should be used, discounting future benefits at a risk-free rate. Second, all funding gaps, regardless of their source, should be closed within five to seven years, which would also eliminate the need for complicated advance funding requirements for seriously underfunded plans, as under present law. Third, benefit increases under flat-benefit plans should be treated either as current-year benefit accruals or as amendments, with the resulting funding gap to be closed promptly; the potential costs of shutdown benefits should also be recognized and funded promptly. Fourth, PBGC guarantees for plan changes, including shutdown benefits, should be phased in over the same period that the corresponding funding gaps are closed. Fifth, any subsidized portion of early-retirement benefits or lump-sum payments should not be made out of pension assets and should not be insured by the PBGC. Sixth, a plan should be frozen whenever the sponsor is delinquent in making payments. Finally, all plan assets should be valued at current market prices.

To immunize DB plan funds from valuation swings, Wilcox reasons that portfolios should be restricted to high-quality debt instruments structured to deliver cash flows as benefit obligations become due. Similarly, for its own portfolio, he would restrict PBGC to investing entirely in zero-coupon bonds structured to mature as its obligations become due. Since these obligations derive from plans already terminated, they are known with certainty up to mortality risk. Wilcox discusses various issues associated with moving to a rational pricing structure for PBGC insurance. He notes that pension analysts have developed models that attempt to price the several risks that exist: aggregate economic risk and sponsor-specific bankruptcy risk as well as risks associated with the pension plan's financing. He suggests that further experience with such models is needed before risk-based premiums are introduced. Wilcox also discusses ways in which the financial positions of DB plans could be made more transparent, and budgetary and accounting requirements could be improved.

Wilcox recognizes that his aggressive approach to DB pension reform may go further than necessary. His main proposed reforms are deliberately redundant in that each reduces the system's present risks in a way that makes the other reforms less necessary. For example, if fund assets are continually marked to market and funding gaps are closed promptly, the risk associated with keeping some equities in a portfolio is reduced. Wilcox nonetheless recommends adoption of all of his reforms on the practical grounds that not all reforms may be adopted simultaneously and that, even if they were, the cost of their redundancy might be very low.