Editors' Summary

THIS ISSUE of the *Brookings Papers on Economic Activity* contains articles, a report, and discussions presented at the fifty-sixth conference of the Brookings Panel on Economic Activity, which was held in Washington, D.C., on April 7 and 8, 1994. The first article examines how changing U.S. trade patterns, particularly trade with developing countries, have affected U.S. manufacturing. The second article presents a detailed empirical study of the role that saving incentive plans, such as 401(k)s and individual retirement accounts, have played in personal and national saving over the past several decades. The third article explores the connection between employer-provided health insurance for retirement. The fourth article looks at recent economic performance in Mexico and the effects of the currency overvaluation that has emerged. The report in this issue examines the reported weakness in the growth of U.S. real wages and compensation over the past 20 years.

GROWING U.S. TRADE with the rest of the world has coincided with three significant trends in U.S. labor markets during the past 15 years: a decline in total manufacturing employment; a widening of the income inequality between high-skilled and low-skilled workers; and especially sharp employment declines in low-skill manufacturing industries. Trade with developing countries grew very fast during this period, rising from 29 percent to 36 percent of U.S. manufactured goods imports between 1978 and 1990. In the first article of this issue, Jeffrey D. Sachs and Howard J. Shatz investigate the link between this growing competition from developing country imports and the developments in U.S. labor markets.

As a starting point for their analysis, Sachs and Shatz review the predictions of what is commonly referred to as the Heckscher-Ohlin-Samuelson model (HOS)—a synthesis of results from classical trade theory. In simple HOS, with two trading countries and two goods, each produced using two factors of production, each country tends to special-

ize in the good that is intensive in the country's abundant factor and trades it with the other. The two factors are skilled and unskilled labor, with the United States relatively abundant in skilled labor and the developing countries relatively abundant in unskilled labor. When barriers fall so as to open up trade, the model predicts the following changes in U.S. markets: a fall in the relative price of less skill-intensive goods; a fall in the relative price of unskilled labor; a rise in the production of more skill-intensive goods and a fall in the production of less skill-intensive goods; a rise in exports of more skill-intensive goods and in imports of less skill-intensive goods; and, within each productive sector, a rise in the ratio of unskilled to skilled labor employed, assuming the overall employment of both types of labor is unaffected.

The authors go on to discuss several amendments to simple HOS that modify its predictions. Allowing for capital as a productive factor, they describe how capital mobility provides additional avenues of influence. If rates of return have not been equalized by free trade and developing countries are relatively short of capital as barriers to foreign investment are reduced, there will be an incentive to relocate U.S. capital there. This increase in the capital available to developing countries would lead to an increase in their production and export of less skill-intensive goods. The imports of these goods into the United States would lead to a fall in the demand for unskilled workers and a fall in their relative wage; a trade deficit would then open and would be paid for by the earnings on the capital invested in the developing countries. This model thus stresses the possibility that opening foreign capital markets, by inducing U.S. capital to move to the developing countries, can reduce U.S. manufacturing employment and the relative wage of unskilled labor, even without the shifts in relative prices and skill intensities that drive the simple HOS model. The authors note that, through this avenue, U.S. multinational corporations may play a distinctive role affecting U.S. trade and wage patterns. As further amendments of the simple HOS framework, they note that technical change may favor the employment of one type of labor over another and that, in a fuller model, the supply of unskilled labor to manufacturing may decline in response to a decline in its wage, either because it can move to the nonmanufacturing sector or for other reasons.

For their empirical work, Sachs and Shatz assemble a new data base that allows them to trace detailed U.S. trade patterns by manufacturing industry and trading partner country. Most of their analysis relies on a disaggregation of trade into 131 manufacturing categories (three-digit SIC classification) and more than 150 trading partners. The authors use production workers as a proxy for unskilled workers and nonproduction workers as a proxy for skilled workers; data on wages and employment for these categories are available for the required detailed industry groups in the Annual Survey of Manufactures and have been shown by others to correlate well with measures of worker skills.

Sachs and Shatz focus on the years 1978–90, a period for which reliable data are available and over which U.S. trade with developing countries expanded significantly. It is also roughly over this interval that the relative wage of unskilled workers fell substantially in the United States and that the ongoing relative decline of manufacturing employment to total employment accelerated. The authors arrange the 131 industries according to skill intensity of U.S. production and aggregate the industries into decile rankings, with decile 1 being the most skill intensive. In the decile 1 industries, which include periodicals and office and computing machines, production workers account for 41 percent of employment; in the decile 10 industries, which include footwear and many kinds of apparel, production workers account for 87 percent of employment. As additional examples of how industries are arrayed in their data, motor vehicles are in decile 8, basic steel is in decile 6, and many kinds of industrial machinery are in decile 3.

The authors use their skill rankings to identify several regularities in 1990 trade patterns. In trade with developing countries, the United States tends to be a net exporter of skill-intensive products and a large net importer of products at the other end of the skill ordering, in accord with predictions of HOS. Similar patterns are evident when trade by industry decile is examined against the wage level of the trading partner, which presumably proxies for the skill intensity of the partner economy. In trade with low-wage partners, the United States has net surpluses in skill-intensive industries and net deficits in non-skill intensive ones. There are no similar regularities for trade with advanced countries, a result consistent with "new trade theory," which predicts most such trade would be in differentiated products within the same industries.

Sachs and Shatz also use their industry data to investigate the pattern of change over time in the net trade deficit of the United States. They calculate the ratio of net trade flows to value-added in each industry, adjusting the value of imports for shipping costs and for the effects of tariffs and other trade restraints in each year so as to express them in values at which they compete in the U.S. market. They show that, between 1978 and 1990, net exports with developing countries fell in all industry skill deciles except decile 2, with the largest declines by far in deciles 9 and 10, in rough accord with their extended HOS model. There is little pattern to the change in trade with developed countries, just as there was little pattern to the 1990 levels of trade. Looking at changes in employment, they show that the employment of unskilled labor fell in all deciles while that of skilled labor rose in deciles 1, 2, and 3 and had little pattern elsewhere. Thus, contrary to the predictions of simple HOS—or of any model that predicts factor intensities will vary with relative factor costs—the skill intensity of production rises in every decile despite the decline in the relative wages of unskilled workers, suggesting other important forces are at work.

The authors extend this analysis with cross-sectoral regressions explaining changes in sectoral employment from 1978 to 1990 by key structural characteristics of sectors: the share of capital in value-added, the share of production workers, and spending on research and development relative to sales. The first two characteristics each had a significant negative effect on employment growth for both production and nonproduction workers, which, the authors argue, is consistent with the idea that international capital mobility and unskilled labor competition from abroad have been important. Research and development spending had a negative effect on production worker employment, which the authors take as evidence that technical change has been associated with a rising ratio of skilled to unskilled workers.

To estimate more formally the relationship between import penetration and employment, Sachs and Shatz start with a counterfactual case in which import penetration as a share of final demand does not rise after 1978, thus requiring more domestic employment than actually occurred. Using 51 manufacturing industries and input-output relations linking final sales to sectoral output and employment requirements, they calculate their counterfactual: what 1990 imports in each sector would have been if the ratio of imports to final demand in each sector had stayed at 1978 levels. They then calculate counterfactual employment totals by industry, divided between production and nonproduction workers. The difference between counterfactual employment and actual employment is their estimate of the job loss from increased net imports. The authors find that shifts in trade with developing countries account for a 6.2 percent decline in production worker employment and a 4.3 percent decline in nonproduction worker employment. By contrast, shifts in trade with developed countries have had almost no net effect on total employment, with a 1.0 percent decline in production worker employment approximately offsetting a 2.2 percent rise in nonproduction worker employment. In total, trade accounts for a net decline of 5.9 percent in manufacturing employment, or 1.2 million jobs, over 1978–90. This is 39 percent of what the authors estimate to be the total shortfall in manufacturing employment over this period.

By skill deciles, only decile 1 industries experienced a substantial employment increase resulting from trade, while decile 10 industries experienced the largest decline resulting from trade, 27 percent. Comparing counterfactual and actual changes in employment by individual industries, the authors find a high correlation between the two, confirming that trade was an important determinant of actual employment changes. However, they also make clear that employment changes in a number of industries, such as the declines in coal and primary metals and the increases in publishing, were largely unrelated to trade.

It is more difficult to assess the effects of trade on relative wages than to assess the effects on employment as discussed thus far. Sachs and Shatz observe that manufacturing accounts for only 16 percent of nonagricultural employment; labor markets are segmented so that industry earnings depend on union premiums and other industry characteristics; and micro-level data are needed to account for detailed worker, as well as industry, characteristics. They add some general observations about why trade effects cannot be expected to explain much of the economywide change in relative wages: to the extent the manufacturing work force is segmented from the rest of the economy, differential demands for skilled and unskilled labor may have contributed importantly to widening wage inequalities within manufacturing but not outside it. On the other hand, to the extent the manufacturing work force is highly integrated with the rest of the economy, even the noticeable declines that they find in the relative demand for unskilled labor in manufacturing would be small relative to the economywide demand for such labor.

Aware of these difficulties, Sachs and Shatz nonetheless look for evidence on the effects of trade with major developing countries on the U.S. wage distribution. Because HOS links relative wage effects to changes in effective prices arising from trade, they calculate "effective prices"—output prices adjusted for total factor productivity—for traded goods sectors. Using import price indexes, they find only a weak and insignificant relation of price changes to the skill mix of employment in the sector. Using domestic price deflators, they find a stronger and significant relation with which they estimate that decile 10 goods fell in price by 9 percent compared with decile 1 goods over 1978–89. But these effects, while significant, are not large enough to account for an important widening of wage differentials throughout the economy. Sachs and Shatz conclude, as have some previous authors, that technical change that substituted capital and skilled labor for unskilled labor must have been a second important development, along with enlarged trade, for recent U.S. wage developments. However, they cannot quantify the importance of such an effect.

The authors more fully examine a third potential development—the possibly special role played by international capital movements in enlarging the trade opportunities of developing countries and depressing wages and employment in competing U.S. industries. They focus, in particular, on the rise in globally integrated production within transnational corporations (TNCs), which results in sales from an affiliate in one country to an affiliate or parent in another country. As an example, they note that in 1990, 28 percent of U.S. exports to Mexico and 34 percent of imports from Mexico were sales between affiliates. They observe that such trade integration, which involves direct investment, brings linkages in research and development, technology, skills, and organizational patterns. And they project that TNC-based trade with developing countries is likely to rise rapidly in the future, encouraged by global and regional agreements such as the Uruguay Round and the North American Free Trade Agreement (NAFTA).

To explore whether TNC-based trade can be expected to increase overall trade by facilitating an international division of labor within firms, they estimate "gravity equations" for exports and imports that predict the volume of trade between the United States and a large number of partner countries as a function of an index of TNC trade and several structural characteristics, such as population, GDP per capita, and geographical distance from the United States. The TNC variable is not significant in 1978 cross-sections. But in 1990 cross-sections, it predicts significantly increased trade volumes: a 10 percentage point increase in trade within TNCs predicts a 40 percent increase in the volume of both overall exports and imports.

Sachs and Shatz go on to explore the importance of TNC trade on the skill structure of overall trade. Within each skill decile, they estimate cross-country regressions explaining net trade by each country's wage relative to the U.S. wage and the TNC index, where the index is allowed to take on a different coefficient for developing and developed economies. For developing countries, they find that higher TNC-based trade is associated with larger U.S. net trade deficits across all skill deciles. No relation is found for developed countries. From their results, the authors surmise that U.S. firms use production facilities in developing countries in two ways, depending on the type of good being produced: as an export platform for reexport to the United States, and as a substitute for U.S. production in supplying the developing country market. They conjecture that the former is most likely for less skill-intensive goods and the latter for more skill-intensive ones, although they acknowledge they do not have the detailed data needed to confirm this conjecture.

Sachs and Shatz assess how continued trade liberalization and increased capital mobility-both of which they expect to intensify-are likely to affect U.S. labor and product markets. They apply their findings to the particular circumstances of China, India, the former socialist economies, and Mexico. In China, they expect continued rapid trade growth to result from rapid growth in Chinese gross domestic product and in foreign direct investment by transnational corporations, which already totaled \$25 billion in 1993. With plausible trends in growth and direct investment, their gravity equation predicts the volume of U.S.-Chinese manufacturing trade, which amounted to 2.4 percent of U.S. manufacturing trade in 1990, will grow by 16 percent a year for the foreseeable future. India is a decade behind China in opening its economy, but it has done so finally, and its GDP is starting to grow. The authors project U.S. trade with India will grow about as fast as with China, but starting from a much lower base. Trade with the former socialist economies has been much smaller than the gravity equation predicted on the basis of income, population, distance, and the like. According to the authors' estimates, if market reforms remove this discrepancy, trade volumes would rise ninefold on that account alone, growing to 6 percent of U.S. trade from less than 1 percent currently. Finally, the authors expect that trade with Mexico will increase as a result of Mexican growth and increased TNC-based trade encouraged by NAFTA. If TNC trade in Mexico became as important as it is now in Canada, they project that that change alone would increase Mexican-U.S. trade by 50 percent.

The projected expansion of developing country trade can be expected to have the same kind of adverse effect on low-skilled workers that the authors have identified from the recent past. They conclude that, while increased trade with the developing countries is in the best overall interests of the United States, the dislocations of low-wage manufacturing workers deserve special attention, such as the Clinton administration's call for enhanced job training and assistance in job relocation.

THE U.S. NATIONAL saving rate, which averaged roughly 8 percent in the 30 years following World War II, fell to about 4.5 percent during the 1980s and has fallen below 2 percent since 1990. While some of this decline reflects increases in government deficits, the personal saving rate has also declined substantially, from an average of 7 percent between 1950 and 1980 to an average of 4.6 percent during the 1990s. These low rates of saving have for some time concerned economists and policymakers, because they imply low rates of capital formation or increasing foreign ownership of U.S. capital, or both. Most recently, these concerns set the political stage for the passage of the budget agreement of 1993, which is intended to reduce government dissaving. They have also motivated proposals to increase personal saving, including cutting capital gains taxes and providing tax-advantaged saving plans. In the second paper of this issue, Eric Engen, William G. Gale, and John Karl Scholz examine the historical evidence on the effectiveness of the tax incentives embodied in individual retirement accounts (IRAs) and 401(k) and Keogh plans. Since 1986, these plans have attracted funds equal to about one-third of total personal saving as measured by the National Income and Product Accounts. The authors focus on the effects of these plans both on personal saving and, taking into account the effects of lost tax revenue, on national saving. Their estimates of these effects are substantially smaller from those of many previous studies.

The authors begin with a description and analysis of the central features of various saving incentive options. The first IRAs were established for workers without pensions in 1974 and provided tax-deductible contributions up to an annual limit, tax-free accrual of earnings, and penalties for early withdrawal of funds. IRAs were modified after 1986, leaving the earnings, but not the contributions, tax free. 401(k) plans, which have many of the same features as the original IRAs, were established in 1978 but are available only to employees of firms that choose to sponsor the plans. Since then, the eligibility, limits, and other features of these plans have been changed several times, but their essential characteristics have remained unchanged.

The authors show how the plans affect the after-tax returns to saving. The value of the tax benefits obviously depends on whether the tax rates an individual faces increase or decrease over time, either because of changes in the individual's tax bracket or because of changes in the tax law. In the simple case where tax brackets do not change and where both contributions and investment earnings are tax deferred, the plans amount to making the income earned on the saving tax free, an important tax preference. A pretax \$100 invested at 8 percent in a taxable account, with a 30 percent marginal tax rate, grows to \$359 after tax after 30 years; over the same period, the same forgone consumption invested in a fully deductible IRA or 401(k) plan grows to \$700. Eliminating the deductibility of contributions, as in present IRAs for higher-income individuals, obviously reduces the attractiveness of the plan and makes the tax advantage important only for long holding periods.

Contributions to 401(k) plans have grown steadily since their inception and are new the dominant form of incentive saving. IRAs grew rapidly following universal eligibility in 1982 but fell dramatically following the restrictions on deductibility under the 1986 tax reform; contributions are now about a quarter of their 1986 value. Keogh and simplified employee pensions, although they grew substantially during the 1980s, are much smaller than the other two plans. Contributions to all saving incentive plans have grown from 0.2 percent of gross domestic product in 1980 to 1.2 percent of GDP by 1990. But whether these plans added to total personal saving, let alone national saving, is a separate matter. The authors show that even though incentive plans provide substantial tax benefits to savers, their theoretical effect on personal saving is ambiguous. First, the income effects of the tax saving unambiguously raise consumption and so reduce saving. These income effects may outweigh the substitution effects of a higher rate of return. Individuals who are saving for retirement, for example, and who have targeted a postretirement level of consumption will be able to meet this target while substantially reducing their saving and enjoying higher levels of consumption during their working years. Second, individuals may save beyond the allowed

limits of the plan. For such individuals, the substitution effect is zero, and the income effect, which remains, reduces saving. This outcome is more likely for IRA accounts with their relatively low limits and for higher-income individuals who are responsible for a disproportionate share of saving.

The authors present an array of empirical evidence aimed at uncovering the effects of tax incentive plans on saving. They first present new evidence on the effects of 401(k) plans using information from the Survey of Income and Program Participation (SIPP). This national household survey provides data on household income and demographics and, among other variables, 401(k) eligibility and participation. For 1987 and 1991, the survey also contains information on 401(k) balances. However, this survey like most others, does not track individual households over periods during which the availability and terms of tax incentive plans vary. Thus, inferences about behavior have to come from comparing the saving of different households, some of which are eligible for and participate in tax incentive plans. This raises two problems. First, participating households have different income and demographic characteristics, which themselves may be responsible for different levels of saving. The authors control for such characteristics when estimating the effect of incentive plans on saving behavior.

Second, and less easy to deal with, households have differences in underlying saving propensities that are not explained by such measurable characteristics. Households with high saving propensities are not only likely to have above-average wealth, controlling for observed characteristics, they are also likely to make up a disproportionate share of households that participate in voluntary plans like IRAs. The authors deal with this second problem by comparing the behavior of two groups: 401(k) participants and IRA participants who are not eligible for 401(k) plans. They reason that households in both these groups have revealed themselves as "savers" and are likely to be more similar in their underlying saving propensities than households selected either randomly or on the basis of eligibility. Since the saving incentives embodied in 401(k) plans have remained essentially unchanged since the early 1980s, whereas the saving incentives behind IRAs were substantially curtailed following 1986, 401(k) participants should have saved relatively more after 1986 if the incentives are important.

The authors run regressions explaining wealth in 1976 and 1991 for

the combined sample of households. The regressions control for time effects and for a number of variables usually thought to influence wealth accumulation, including family income, age, education, defined benefit pension coverage, and various occupational and demographic variables. The estimated coefficients for these variables by and large accord with life-cycle predictions for several measures of wealth. The crucial variable in the regression is an interaction dummy variable, identifying 1991 401(k) participants. This variable should be positive and significant if 401(k) plans raise private saving. Instead, for a variety of specifications, it is insignificant and, in all cases but one, negative.

Although they prefer their own comparison groups because they use households that are likely to have similar underlying saving propensities, the authors also examine an alternative comparison group that is used by James M. Poterba, Steven F. Venti, and David A. Wise in their study of the effects of 401(k) plans: families without 401(k)s and families with 401(k)s but without IRAs. Regressions controlling for the same economic and demographic variables as above yield insignificant coefficients, but the point estimates indicate that 401(k) participating households increased saving by approximately \$400 more than nonparticipants between 1987 and 1991, implying that between 10 and 14 percent of 401(k) contributions represent a net increase in national saving. The authors observe that, since low-saving households are more likely to be in the non-401(k) population and high-saving households are more likely to be in the 401(k) population, these estimates are likely to suffer from selection bias. They demonstrate the likely importance of this bias with regressions showing that eligibility is positively and significantly related to median financial assets and net worth, both measured excluding 401(k) assets; its relation to net financial assets is negative and significant only in 1987. Since 401(k) assets of eligibles should reduce their non-401(k) asset holdings, these results, if anything, probably understate the positive correlation between saving tastes and eligibility and understate the bias in comparing 401(k) participants and nonparticipants.

As already noted, one source of uncertainty about the results discussed thus far is that the data come from surveys that do not track the same households through time. The authors examine alternative data from the 1979–88 IRS Michigan Tax Panel, which follow the same households over the period during which statutory limits on IRAs changed. Although the data set has this major advantage, it has other disadvantages: asset balances are not available and have to be inferred from reported interest and dividend receipts; the returns provide only limited information on household characteristics; and there is nonrandom attrition of households from the panel. A simple comparison of the saving of continuing IRA contributors with the saving of contributors made newly eligible in 1982 shows that continuing contributors increased their saving more than new contributors. Hence the comparison provides no support for the proposition that expanding IRA eligibility stimulated saving. By contrast, the estimates suggest that an increase in the contribution limit—which differs for single and married taxpayers, for single and two-worker households, and, after the 1986 reform, by income-has a positive effect on saving; in the most favorable regression, gross private saving increases by 31 percent of an increase in the limit. Taking into account the estimates of the forgone tax revenue, this implies that about 4 percent of IRA contributions represented an increase in national saving during 1982-86. However, the authors find these results are not very robust; when they drop the last year, thus eliminating the period when high-income households had deductibility restricted, the estimates become of erratic sign and insignificant.

None of these other examinations provide results strong enough to alter the authors' main conclusions that saving incentive plans have not raised private wealth. How can this conclusion be reconciled with the very substantial increase in contributions to these plans since their inception? The authors note that earlier research showed some contributions have substituted for other private pension plans; the extent of substitution between 401(k)s and IRAs provides additional evidence on the substitutability of different types of saving. If households regard IRAs and 401(k)s as substitutes, participation in IRAs after 1986, when they became less attractive and more restricted, should have declined more among 401(k) eligibles than noneligibles. The authors' estimated substitution into 401(k)s from IRAs is statistically significant and accounts for approximately half of the increase in 401(k) balances of affected households. The estimates from the SIPP data also provide evidence that defined benefit pension plans and 401(k) balances are substitutes; having a defined benefit plan reduces expected 401(k) balances by about \$3,000.

The empirical estimates of the effects of IRAs and 401(k)s on personal and national saving are necessarily short-run effects, given the relatively

limited time they have been in existence. The long-run consequences may be different since the typical household will face improved saving opportunities throughout its life. In order to analyze the long-run effects of improved saving incentives and the complicated transition to a longrun steady state, the authors simulate a life-cycle model that allows for stochastic earnings and uncertain lifespan and that incorporates the salient features of the current tax law and saving incentive plans, including limits on annual contributions and penalties on early withdrawals. They rely on previous studies to provide realistic estimates of the life-cycle earnings profiles for three different education classes. The model does a reasonable job of matching empirical patterns of consumption, saving, and wealth, including the mean and median levels of asset holdings, the percentage of households that hold saving incentive plans, and the percentage of contributors that are at the statutory limits. It is only at striking variance with the data in its underrepresentation of households with zero wealth.

Simulations of the model show that positive long-run effects on saving and wealth are quite modest and do not arrive for decades. In the long run, introducing IRAs increases the level of national saving by about 3 to 5 percent, with the asset-to-income ratio rising from 3.6 to 3.7 or 3.8 depending on the contribution limit. The 401(k) plans, with their higher contribution limit, are somewhat more effective, especially when employers make matching contributions. In the long run, they increase saving by between 8 and 17 percent, depending on the particular features of the plan. The authors conclude that both IRAs and 401(k) plans can stimulate national saving in the long run but that the magnitude of the increase is small relative to recent declines.

The simulated short-run effects of introducing saving incentives are quite different from the long-run effects, and the authors regard them as broadly consistent with their other empirical results. When incentives are introduced, the typical household faces improved opportunities for only a portion of its lifetime, and most will have already accumulated assets through past saving, complicating the response that the life-cycle theory would predict. Households with low levels of pre-IRA saving will face unaltered marginal rates of return until they have moved accumulated savings into IRAs. Until that time, their consumption should rise rather than fall. After the portfolio has been shifted, the marginal rate of return may increase, and consumption may fall. The proportion of households with savings that predate the reform will gradually fall, and, as they do, aggregate saving will more closely resemble the long-run equilibrium. In the authors' IRA simulations, the ratio of private assets to income falls for more than 15 years and does not return to its pre-IRA value for 36 years. In the transition period, national saving falls by much more than private saving because of the revenue losses that come from the deductibility of IRA contributions.

Taken as a whole the authors' results suggest that a large part of the saving that has flowed into tax incentive plans was by "high savers," individuals who would be doing above-average saving even in the absence of such plans and for whom the tax-sheltered savings are largely substitutes for other forms of saving. Earlier studies, which suggest that the tax incentive plans provide an effective tool for increasing national saving, are likely to be much too optimistic.

A MAJORITY of U.S. citizens believe that everyone should have guaranteed access to a reasonable level of medical care independent of their employment and economic status. This belief, together with the view that the current system is an inefficient and costly way of providing medical services, provides the main impetus for reform. Economists are also concerned that individuals are locked into their current jobs under the present, largely employer-based system, to the disadvantage of individuals and, possibly, overall economic efficiency. And, just as employerprovided health insurance is thought to deter job mobility, it is also likely to deter retirement, with some of the same implications. In the third paper of this issue, Brigitte Madrian examines the importance of this phenomenon for workers' retirement decisions and the implications of universal postretirement health insurance coverage for labor force participation.

As a prelude to her analysis, Madrian presents a variety of reasons why individuals may regard health insurance coverage as important when making retirement decisions. The frequency and severity of health problems increase dramatically as individuals approach retirement age. For example, she notes that more than 20 percent of individuals in the 55–64 age group suffer from three or more serious health problems, such as stroke, cancer, arthritis, or heart disease. This is roughly double the percent of the 45–54 age group who are similarly afflicted. Average annual medical expenditures, conditional on their being greater than zero, rise from \$3,106 for an individual aged 45–54, to \$4,125 for an individual aged 55–64, to \$6,348 for an individual aged 65–74. Among individuals aged 55–64, almost one-third reported being in fair or poor health. Given the increased frequency and severity of health problems as individuals approach retirement age, it is not surprising that survey results report that 63 percent of working Americans would delay retirement until becoming eligible for Medicare at age 65 if their employers were not going to provide them with health coverage.

Despite a vast literature on the determinants of retirement, very little work has been done on the connection between health insurance and retirement. The availability of appropriate data has been a major obstacle to such research. Although firms began providing postretirement health insurance benefits in the late 1940s, and Medicare's assumption of a major share of health care expenditures for the elderly in 1965 further stimulated their provision, there are no consistent time series data to show the rise in the availability of postretirement health insurance. Cross-section data have also been limited; most previous studies of the effect of health insurance on retirement have depended on unrepresentative samples or have had to impute values for important variables. Madrian exploits three previously unutilized data sets, the 1987 National Medical Expenditure Survey (NMES) and two supplements of the Survey of Income and Program Participation (SIPP). All of these data sets provide a representative sample of retired individuals and include information on whether they had access to employer-provided retiree health insurance and their age at retirement.

Madrian provides a picture of the structure of postretirement insurance benefits using simple tabulations of the cross-sectional data and information from a variety of other sources. The incidence of postretirement benefits by age groups and year of retirement shows the substantial growth in coverage that has taken place in the postwar period. For all of Madrian's data sets, the fraction of retired employees who received some form of postretirement health insurance from their employers grew steadily from less than one-third in the early 1960s to nearly onehalf by the late 1980s. A similar increase in coverage is evident when the data are viewed by age group. For example, 45 percent of individuals in the 65–69 age group have employer-provided coverage whereas only slightly more than a quarter of individuals in the 80–84 age group have it. There are substantial and systematic differences by firm size and industry in the provision of postretirement health insurance coverage. Only 15 percent of firms with fewer than 100 workers provide such benefits while more than 80 percent of firms with more than 2,500 employees provide them. More than two-thirds of individuals retiring from jobs in public administration receive postretirement health insurance, but less than 15 percent of those in sales, personal services, and agriculture receive it.

Madrian shows that industries that are more likely to provide pensions and health insurance to active employees are more likely to provide postretirement health insurance as well. But surprisingly, and fortunately for the later econometric work, eligibility for postretirement coverage is not necessarily tied to the receipt of other retirement benefits; for only about one-quarter of workers are postretirement health benefits contingent on having a pension, and for more than one-third of employees there are no prerequisites. Madrian also documents that the costs to retirees of employer-provided health coverage bear little resemblance to the cost to employers. The coverage is typically identical to that provided to active employees and the contributions by retirees are relatively modest; in 1990, the median contribution for family coverage was \$55 for pre-65 retirees and \$31 for post-65 retirees. The costs to employers of providing retiree health coverage differ radically with the retiree's age; because Medicare pays for many expenditures of retirees over 65, their costs to employers are typically half the costs for retirees under 65.

The centerpiece of Madrian's estimation of the effect of health insurance on retirement is a series of regressions relating retirement age to retirement health coverage. Because no longitudinal data are currently available with information about whether health insurance coverage of working individuals will continue after retirement, she makes use of the information on those who have already retired. Her data sets, the NMES and two subsets of SIPP panel data, all provide information about age at time of retirement and retirement health coverage but differ in the other information they provide. She restricts her analysis to men because at the time that those in her retrospective sample retired, the labor force participation of women was much lower than it is today and because the retirement decision for men seems better defined.

Madrian is aware that her estimation strategy has censoring problems. Most important, she is forced to use a truncated sample consisting

xxiv

of retired individuals, since she does not observe individuals who have not yet retired. In a univariate regression where only the availability of postretirement health insurance is used to explain retirement, exclusion of individuals who have not yet retired will lead to an underestimate of the effect of insurance coverage on retirement age. In multivariate regressions, the direction of the bias is ambiguous. Madrian shows that at younger ages, the truncation of the sample is substantial, with only 24 percent of those aged 55–59 reporting ever having retired; at age 70, however, 85 percent of individuals report being out of the labor force, so there are few nonretired left to bias the results. Madrian deals with this truncation problem in two ways: by running separate regressions for the populations of individuals over 54, over 64, and over 70, and by using a truncated regression model, which explicitly adjusts for this selection bias.

A second censoring problem arises because some individuals have died since retirement and are not in the sample. While raising the age cutoff of the sample decreases the first type of censoring problem, it obviously increases this second potential source of bias. Since individuals with poor health tend to retire early, a bias will occur if individuals with poor life expectancy are either more or less likely to have retiree health coverage. Since there is some evidence both that income and mortality are negatively correlated and that retiree health coverage and income are likely to be positively correlated, there may be an upward bias in the coefficient on health coverage because individuals who have died are not in the sample. Madrian notes that most mortality occurs after age 65, when all individuals have coverage from Medicare, so that this bias may be small. Madrian also notes that another bias arises if jobs are selected because they provide health insurance for early retirement, but she argues that this is likely to be unimportant because most individuals do not know whether they will receive such coverage when they take a job.

Although they differ in the precise definition of retirement and in the other information available about individuals and their preretirement employment, the three sets of data give broadly consistent results. Retiree health insurance coverage is always significant in explaining retirement age. Restricting the samples to older individuals, or using her truncated regression model to reduce the selection bias arising from the missing nonretirees, systematically increases the estimated effect of postretirement health coverage. The estimated effects from the NMES

sample are largest. Controlling for education and race with a complete set of cohort dummies, the regressions suggest that coverage decreases retirement age by between a year and a year and a half. The two subsets of the SIPP data provide somewhat smaller estimates, with the truncated regression model showing an effect of postretirement coverage ranging from a 0.4-year to a 1.2-year reduction in retirement age.

The SIPP data include pension information that permits some further tests. Restricting the sample to those with pension income, thereby diminishing the chances that retirement health coverage is proxying for early retirement pension incentives, decreases the estimate, but not significantly or by an important amount. Madrian recognizes that the lack of precise information on the correlation between pension incentives for early retirement and retirement health coverage leaves one potential source of bias unresolved. She utilizes the SIPP subsample on job characteristics to explore the relationship between retirement health coverage and pensions with inconclusive results and suggests that more work needs to be done on disentangling the effect of these two possible incentives.

Postretirement coverage might be expected to be most important for individuals considering early retirement before Medicare becomes available. To investigate this, Madrian estimates a probit equation predicting whether an individual retires before age 65, using the same basic variables as in the earlier regressions. For the NMES sample, postretirement health insurance is estimated to increase the probability of retiring before age 65 by 15 percentage points; for the SIPP samples, the estimated increase is about half as large. Madrian's overall conclusion is that retiree health insurance is a strong predictor of early retirement; her subjective estimate, considering all the results, is that the availability of such coverage decreases retirement age by about one year.

Madrian's estimates suggest that the increase in postretirement health insurance coverage may be an important contributor to the declining labor force participation of older workers during the past few decades. In 1960, 34 percent of men over 65 were working; by 1990, this percentage had fallen to 16 percent. Previous authors have found that between one-third and one-half of this decline in labor force participation of older individuals is explained by improvements in private pensions and social security. Making use of the information in her samples on the increased availability of retiree health insurance and the probit results, Madrian suggests that the increased availability of retiree health insurance can account for 10 to 21 percent of the aggregate reduction in participation, about half the effect estimated to arise from private pensions.

Madrian's results also suggest how national health care reform is likely to affect labor force participation. Using a subjective average of her estimates. Madrian calculates that a move to universal health care would reduce the labor force participation of the most affected group, men aged 60–64, by more than 4 percentage points. She estimates the overall reduction in employment for men aged 50-69 would be about 450,000. Madrian recognizes that the effects of national health care reform on employment depend on the response of the demand for workers, as well as the response of their supply. Firms may respond to mandatory and universal coverage in a variety of ways. Under some proposed plans, community-rating will make older workers relatively less expensive and may increase the demand for their services; firms not now providing health coverage may decrease their demand; and firms may alter other features of their compensation packages to offset the incentive effects of mandated health coverage. Madrian also stresses that specific features of retiree health insurance make a substantial difference to the incentive for early retirement. While she believes that her results establish the importance of retiree health insurance to labor force participation, she concludes that much more research remains to be done before the effect of health care reform on the labor market can be fully assessed.

BY EARLY in this decade, Mexico had overcome severe problems of foreign debt service and inflation and was coming to be viewed as a model of economic reform and development. The Brady plan, instituted in 1990 to improve the credit quality of Mexico's debt and reduce its interest burden, was the final step in providing financial stability and in making Mexico attractive to foreign investors. Yet, in the past few years, growth of output and employment has slowed, at least in the formal sector of the economy for which data are available; political stability has been brought into question by the uprising in Chiapas and the assassination of Luis Donaldo Colosio, the presidential candidate of the ruling party; and the peso has come under pressure in foreign exchange markets, driving Mexican interest rates to levels that threaten to damage the balance sheets of firms and push the economy into recession. In the fourth paper of this volume, Rudiger Dornbusch and Alejandro Werner diagnose Mexico's current problems and propose dealing with them through a substantial currency devaluation that is linked to a new incomes policy.

The authors briefly recount the economic history that led to Mexico's present situation. From the mid-1950s until the first oil shock in 1973, Mexico enjoyed an extended period of economic growth and relative price stability, with per capita gross domestic product rising 3.1 percent a year and consumer price inflation averaging 3.9 percent a year. The oil price increases of the 1970s, which brought Mexico dramatic increases in export revenues, led to highly expansionary budget policies, large increases in foreign debt, and an overvalued exchange rate. When world interest rates soared and oil prices fell, Mexico, in the summer of 1982, became unable to service its foreign debt and its economic policies came to be dominated by the desire to keep capital from fleeing the country and the need to meet interest payments. Ultimately, lower world interest rates and the Brady plan's debt restructuring brought the foreign debt service under control. But the authors regard important structural reforms of this period as equally crucial, both to restoring confidence in Mexico's finances and to setting the domestic economy back on a noninflationary growth path.

Budget deficit reduction, privatization, and trade liberalization were the key structural reforms undertaken in the wake of the debt crisis. Between 1982 and 1992, expenditure cuts and tax increases both contributed to reducing the operational budget deficit by 9 percent of GDP, resulting in a 3.6 percent surplus by 1992. Over the same period, privatization cut the number of state enterprises from 1,155 to 217 and the number of workers employed in state enterprises by 40 percent. Privatization aided efficiency directly, contributed to deficit reduction, and, by signaling a change in the government's posture toward free enterprise, helped make Mexico more attractive to foreign investors. Import quotas, which had been pervasive, were virtually eliminated and tariffs were reduced to an average level of 11 percent. All remaining barriers to free trade in goods and services within North America are scheduled to fall as the provisions in the North American Free Trade Agreement (NAFTA) are implemented.

By the late 1980s, two other key policies were introduced to control

xxviii

inflation, which had reached a rate of nearly 200 percent a year. The government, in concert with labor and business leaders, adopted a series of *pactos*, a form of incomes policy aimed at controlling wages and prices. And, as part of its contribution to the *pacto*, the government undertook to control the nominal peso-dollar exchange rate, alternately fixing it or keeping it on a gradually increasing path. While Dornbusch and Werner credit these policies with bringing inflation under control, they also point out that they substantially appreciated the real exchange rate, a key development that they examine critically using two alternative models.

The first model, which appears to be guiding Mexican policymakers, is essentially neoclassical and assumes that labor and exchange markets stay near equilibrium. The authors consider the equilibrium model's predictions of how the economy would respond to three policy changes of recent years. The first is a reduction of the budget deficit, which, without access to international capital markets, reduces demand for home goods, lowers the cost of capital, increases investment in the traded goods sector, and depreciates the real exchange rate. The second is trade liberalization, which also reduces the demand for home goods and requires a real depreciation to sustain employment levels. The third is opening financial markets to international capital (equivalent to reducing the risk premium confronting foreign investors), which increases asset prices and leads to capital flows into the country, thus appreciating the real exchange rate, encouraging capital formation, and eventually raising real wages.

In light of these predicted reactions, believers in the equilibrium view see no problem with the appreciation of the real exchange rate that followed from the increasing attractiveness of Mexican financial markets to foreign capital. According to this view, the large increase in the current account deficit in recent years is the counterpart to desirable capital inflows, and the weakness in employment and output is a transitional feature of the economy's response to reform and restructuring.

Dornbusch and Werner believe that an alternative model, which stresses the possibility of disequilibrium resulting from the operation of the *pactos* with their exchange rate targets, provides a more realistic interpretation of recent events. In this alternative model, which the authors illustrate with a simulation, price inflation normally depends on wage inflation and the change in the exchange rate. Wage inflation, in turn, normally depends on past price inflation and the deviation of output from potential output. But *pactos* can intercede to fix wages, prices, and the exchange rate at policy-determined rates of growth. The model normally has a high degree of inflation persistence; but this process is interrupted when the nominal exchange rate is fixed, allowing a rapid disinflation to take place. As a crucial side effect of this interruption, the real exchange rate appreciates sharply and the current account moves toward deficit. The return to a competitive real exchange rate takes place only gradually and only through a sustained depression of output and employment.

In support of their disequilibrium view, Dornbusch and Werner offer several observations about recent trends. They show that the structure of imports has moved toward consumer goods, which especially benefited from trade liberalization, rather than toward capital goods, as might have been expected if capital formation were driving the growing deficit in the current account, as in the equilibrium model. In 1992, capital goods accounted for only 19 percent of total imports, with intermediate goods accounting for 69 percent. They also note that the disequilibrium model explains the sustained weakness in output and employment while the equilibrium approach dismisses them as only transitional. And, using regressions, they provide evidence that recent exchange rate movements are not part of an equilibrium process. The regressions seek to explain financial markets' expectations of devaluation, which they measure by the differential between peso-denominated and dollardenominated short-term securities. Using data spanning April 1988-June 1993, they find evidence that these expectations depend positively on the trade deficit. They argue that if markets perceived the growing trade deficit as an equilibrium phenomenon, no such relation would be found.

Having concluded that the peso has become overvalued and that this overvaluation, together with the high real interest rates needed to sustain it, represents a disequilibrium that is crippling to employment and output and that interferes with the growth of exports, the authors offer a strong policy prescription. They argue for a 20 percent devaluation as part of a new *pacto* that would also moderate wage and price increases so as to prevent the devaluation from sharply worsening the inflation rate. They believe that the time for such actions is now, in part because they judge that support for a new *pacto* would be forthcoming today but perhaps not after this fall's election. If *pacto* support is not forthcoming,

they suggest using the present crisis to discard the *pacto* process and move to a floating exchange rate with lower interest rates.

STATISTICS ABOUT changes in U.S. labor income have fueled concerns that real wages are not keeping up with productivity growth and, by some measures, are actually declining; that sharply rising fringe benefit costs are undermining gains in take-home pay; and that workers' pay is rising faster in other advanced countries than in the United States. In a report in this volume, Barry Bosworth and George L. Perry examine these concerns by looking critically at the price and wage indexes that underlie them and clarifying how real wage developments have been related to productivity.

Over long periods, growth in real hourly compensation has generally equaled growth in productivity for the private nonfarm economy. After 1973, productivity growth slowed markedly, but real compensation slowed even more, so that, by 1993, it had fallen behind productivity by 10 percent. Many interpret this to mean that labor has not been receiving its historic share of productivity gains. But the authors show that this is not the case. They divide the divergence between the growth in real compensation and productivity into two parts-a sharp fall in labor's terms of trade, meaning that the prices workers pay for what they buy have been rising faster than the prices of what they produce; and a small decline in real compensation relative to productivity when compensation is deflated by the prices of what labor produces. The second component indicates how productivity gains in the private nonfarm sector are divided, and the shortfall in compensation deflated in this way-real compensation gains trailing productivity gains by 0.1 percent a year—is not unusual. The first component, the decline in labor's terms of trade, is unusual and considered in some detail by the authors.

Bosworth and Perry show that the price index used by the Bureau of Labor Statistics for measuring what workers buy—the consumer price index (CPI) for urban workers—is itself biased. They point out that a major change in the way homeownership costs were measured after 1982 severely overstates the rise in the CPI for periods spanning the change. These costs used to be measured, in large part, by mortgage interest rates; since 1983, they have been measured by rental equivalence. Because the changeover occurred near a historic peak for interest rates, the CPI was pushed up by the rise in rates using the old method but was not pushed back down by the subsequent decline in rates after the methodology was changed. This upward bias in changes in the CPI created a downward bias in the official measure of real compensation growth. The authors show that the personal consumption expenditure (PCE) deflator used in the national income accounts, which is an alternative index of consumer prices and which has been measured on a consistent basis throughout the period, rose 5.5 percent less than the CPI over 1973–93.

Even with the PCE deflator, substantial variations in labor's terms of trade remain, with an improvement of 3 percent over 1973–83 followed by an unprecedented 6 percent worsening over the past decade. The authors find that neither the relative price of imports nor the relative price of farm products—two usual suspects—explain this worsening. However, they find about half the worsening can be accounted for by two other developments: the decline in the relative price of computers, which are much more important in the output deflator than in the consumption deflator, and the increase in the relative cost of owner-occupied housing, which is largely excluded from nonfarm business output and its deflator.

Bosworth and Perry note that the most dismal picture of workers' real wages is provided by the index of real hourly earnings for nonsupervisory workers, which is published monthly by the Bureau of Labor Statistics. In the 20 years between 1973 and 1993, real earnings by this measure have declined a startling 15 percent. The distortions in the CPI described above contribute to this decline, since the real earnings index is formed using the CPI to deflate the nominal hourly earnings series. But the authors show that a bigger problem lies with the nominal hourly earnings index itself, which has risen much less than either of two alternative measures-hourly wages in the employment cost index (ECI) and hourly wages as measured in the national accounts data (which is calculated by the authors from published data). They explain that the earnings index has very likely developed a persistent bias because it is based on sampling techniques that are not properly benchmarked to account for the changing universe of firms and possibly because firms do not properly distinguish between nonsupervisory and other workers.

Comparing hourly compensation as measured in the national accounts data and as measured by the ECI, Bosworth and Perry show that both tell the same story over the entire 1980–93 interval for which ECI data are available: real compensation, using the PCE deflator, rose by

xxxii

about 0.5 percent a year. But, over shorter intervals, the two measures have diverged. They point out differences in the methodologies behind the two, the most important of which is probably that the national accounts estimates are affected by shifts in the mix of high- and low-wage jobs in the economy while the ECI adjusts for such shifts. On these grounds, the ECI more closely reflects the compensation costs confronting individual firms, while the national accounts estimate captures, to some extent, changes in job opportunities confronting individual workers.

The authors also show that fringe benefit costs—the difference between compensation and wages—which are commonly believed to be driving a growing wedge between labor costs to employers and takehome pay of workers, have in fact been nearly stable for years. After growing from 8 percent of total compensation in 1960 to 16 percent in 1979, the cost of fringe benefits moved up only one point further, to 17 percent of compensation by 1993. The main reason for this stability is that the rise in payments for health insurance has been offset by a decline in employer contributions to private pension programs. That decline, in turn, resulted from the high returns on pension portfolios and a shift to defined contribution programs, primarily 401(k) plans.

Bosworth and Perry turn to a broad comparison of U.S. compensation trends with those in six other industrial nations over 1979–92. Variations in labor's terms of trade—the difference between output deflators and consumption deflators—make some difference in these comparisons over intervals when the dollar exchange rate was changing sharply. U.S. real compensation trailed the others' real compensation by less when using PCE deflators than when using GDP deflators during 1979–85, when the dollar appreciated sharply. The reverse was true when the dollar subsequently declined. Using GDP deflators, which the authors consider most relevant for international comparisons, the growth in U.S. real compensation per worker for the private sector trailed the average growth of real compensation in other countries by 1.0 percent a year over the entire period, with little difference between the two subperiods. Slower growth in U.S. output per worker fully accounts for the slower growth in U.S. real compensation.

Bosworth and Perry point out that the slower relative growth of U.S. productivity and real compensation reveals only one part of the relative performance of labor markets. They show that U.S. job creation has

been far superior to that in the other countries, except that in Japan, which has done about as well. Comparing the United States with the European countries in their sample—Great Britain, France, Italy, and Germany—they calculate that, if U.S. performance had just matched the European average over 1979–92, the U.S. unemployment rate would have been 9.9 percent in 1992 rather than 7.4 percent and U.S. employment that year would have been 9.6 million lower than it was; 6 million discouraged workers would not have been counted as unemployed because they would have stopped actively looking for jobs.

xxxiv