Summary of the Papers

THE FIFTH meeting of the Brookings Panel on Microeconomics met in Washington in December of 1991, and the papers covered a range of issues in technology, productivity, labor markets, financial management, and the cost of a carbon tax. Joseph Farrell and Carl Shapiro examined the Federal Communications Commission's decision to set a standard for high-definition television. Ann Friedlaender, Ernst Berndt, and Gerard McCullough looked at how management characteristics affect productivity in the railroad industry. Martin Neil Baily, Charles Hulten, and David Campbell examined individual manufacturing plants and asked why some are more productive than others. Michael Boozer, Alan Krueger, and Shari Wolkon examined differences in school quality for black and white students since the landmark desegregation decision *Brown* v. *Board of Education* in 1954. They looked particularly at the gap between white and black students in computer use and explored whether that gap translates into differences between black wages and white wages. Josef Lakonishok, Andrei Shleifer, and Robert Vishny studied the organization of the pension management industry and why the industry continues to exist even though it provides negative value. In the final paper in this issue, Dale Jorgenson, Daniel Slesnick, and Peter Wilcoxen examined the potential effect of a carbon tax on the economy—an important issue in the debate on global warming.

Farrell and Shapiro on High-Definition TV

The technology is now available to provide high-definition television (HDTV) that will offer viewers better sound and a much clearer and larger picture with the same shape as a movie screen. The FCC has decided not to accept the technical standards and specifications that

have been developed in other countries, but rather to solicit alternative proposed HDTV systems and then select a unique set of standards for the United States. This decision is seen by some observers as creating an opportunity for U.S. companies to reestablish themselves in the production of television sets and related equipment. Other observers, however, see a danger that the FCC will saddle viewers with an inferior system, while the manufacture of the new equipment will still take place primarily overseas.

In this paper Joseph Farrell and Carl Shapiro examine the economics of standard setting for HDTV and the subsequent adoption of the chosen technology. The authors look first at the alternative procedures that could be followed in selecting a standard, and they point out that, in principle, this could be done by means of a market rivalry with no selection process by the FCC. In this case several different systems would coexist for a time before one became dominant. The VHS and Beta formats for VCRs are an example of such a process. A second alternative procedure that could have been followed in the United States is to use a very centralized decision process, as the Japanese and Europeans have done, with a standard selected early by a governmental body and then promoted by it.

In practice, the FCC has encouraged parallel development of alternative systems by means of a competition among different consortia. The FCC will then choose which system will be adopted from among the alternatives. The actual procedure allows more competition than the Japanese or Europeans allowed, but it stops short of a full market rivalry.

In setting up the competition among the different systems, the FCC did not encourage an unrestricted search of the possible technological strategies. The commission made clear that it wished to protect the existing terrestrial broadcasters, and this has meant that HDTV signals are to be required to fit in the same bandwidth as current TV signals.

After a system has been selected, there is the important question of adoption. And if the system that is selected is one that leads to high-cost sets, this will discourage consumers from purchasing them. This is particularly true because there are cheaper ways of improving the quality of the current TV picture. As these become available they may undermine the demand for the expensive HDTV sets that appear on the market.

Farrell and Shapiro turn next to a historical review of the development of HDTV. The Japanese were the pioneers in this technology, with experimental transmissions as early as 1979. The Japanese government coordinated and subsidized the development of the equipment needed for studio production of programming, transmission of the signals, and reception of the improved pictures. The Japanese developed "Hi-Vision," a system with roughly twice as many lines as the traditional signal, and they began regular broadcasts with HiVision in 1989.

The Japanese chose not to make their HDTV system directly compatible with the old system. The old sets will not be able to receive the new transmissions, and the new sets will incorporate a converter to allow them to receive the old signals. And the Japanese made the decision not to confine the new signals to the same bandwidth as the old system. Instead they designed a system for direct satellite transmission.

The Europeans started an HDTV joint venture in 1986 with $180 million of funding from the European Community (EC). More has been spent since then by the participating European electronics companies. The EC has supported a satellite system and encouraged satellite broadcasters to adopt an HDTV standard, the "D2-MAC" system. But the broadcasters are resisting this system, arguing that it gives little advantage over enhanced versions of the existing European television system and that the sets will be too expensive to allow widespread adoption.

Turning to the development of HDTV in the United States, the authors argue that it is unfortunate that the broadcasters have so much influence with the FCC. This has slowed the process of development of HDTV because the broadcasters did not see any significant gain to them in revenues from HDTV. And it has meant that bandwidth restrictions have been adopted that are detrimental to the ultimate picture quality. The FCC proposals have effectively foreclosed the possibility of a satellite system that is incompatible with terrestrial broadcasting. A major cost of the FCC's decisions, say the authors, is that scarce electromagnetic spectrum cannot now be used for other purposes. The authors also fault the FCC for demanding a full HDTV system rather than considering more seriously the ways of enhancing the existing signal. The FCC, they say, has opted for a high-cost, gold-plated approach because it has not paid enough attention to manufacturing costs.

Farrell and Shapiro turn next to the general issue of technology choice when there are "network externalities." For most products or services, different companies can choose their own technological approach, and the market can decide which one to choose. HDTV would be much more costly, however, with a variety of incompatible systems. The market cannot, they argue, solve the coordination problem involved in picking a standard. And a monopoly system might end up emerging. Farrell and Shapiro note also that the industry advisory committee to the FCC is helping to avoid the problems that occur when the choice of an HDTV standard is made without the market.

Given the use of a governmental body to set standards, how should it operate? The authors examine static and dynamic models of the choice between quality and cost, and they draw the following implication for the FCC from these. The FCC should choose a system that appeals to the marginal user and not select the very high level of quality that only the most extreme technology buffs will buy. If many consumers buy the new sets quickly, then this will encourage more programming using the new system, and the adoption process will be speeded up.

Farrell and Shapiro then look at the auction mechanism for choosing a standard. Under such a mechanism, proponents of the alternative HDTV systems would bid for the right to have their system adopted. Would this be desirable? They note that an auction system may lead optimistic bidders to win, even if their HDTV systems are not the best. The authors' analysis also focuses on royalty policy. They find that requiring the winning technology to be licensed at reasonable cost does limit monopoly power, but at the cost of exacerbating the problem of overoptimistic bidders.

To provide some empirical insights into the pattern of consumer technology adoption, the authors examine the introduction of color television, cellular telephones, compact disks, video games, and facsimile machines. In some of these cases, the adoption of the new technologies was very slow—color television and fax machines, for example—while in others it was more rapid. Fax machines languished for many years with very low usage until the technology improved and the cost came down. Then the technology took off. The authors argue that these case studies reinforce their concern that the FCC choose a standard for HDTV that results in rapid adoption.

The authors conclude with a look at prospects for the adoption of HDTV. It is hard to predict the speed of adoption, they say, but they note that Americans watch so much television that even a modest increase in the pleasure of viewing would make an investment in HDTV worthwhile.

Friedlaender, Berndt, and McCullough on Railroad Management

Since the passage of the Staggers Act in 1980, which largely deregulated U.S. railroads, the industry has undergone major changes. The number of Class I railroads has fallen to 14 from 37, employment has fallen 52 percent, and the route mileage has fallen 29 percent. Many economists believe that regulation had encouraged inefficiencies in management. In this paper Ann Friedlaender, Ernst Berndt, and Gerard McCullough explore how regulation has affected managerial effectiveness and the relation of this to operating efficiency.

Under regulation, the Interstate Commerce Commission (ICC) and the Congress were the primary determinants of the railroad industry's behavior. Congress placed a high value on serving a broad geographic base, and the ICC refused to allow the abandonment of uneconomic lines. Rail labor was a major beneficiary of regulation as workers received above-market wages. Rail management, largely insulated from competitive pressures, was able to live a quiet life, and some shippers were favored in a rate structure that made some customers subsidize others.

Following deregulation, stockholders became more important, and managers moved to raise profitability by improving efficiency and increasing rates. Despite these changes, rates of return in the railroad industry have remained low. Since 1984 competition has held down rates, with real rates stable or falling. As a consequence, railroad companies have been investing in nonrailroad activities. The industry has also been affected by a wave of takeovers and mergers, and this has been an important reason why managers have found ways to increase stockholder value.

The authors turn next to an assessment of the managerial changes

that have taken place in the Burlington Northern Railroad Company. (Other case studies are in an appendix to this paper.) They note that, following deregulation, a nonrailroad chief executive was brought in to run the company. The company undertook a program of diversification and also took steps to discourage takeovers, including increases in its indebtedness. In 1990 labor productivity in railroad operations was more than double what it had been in 1982.

To go beyond the lessons from this and other case studies, the authors undertake an econometric study of the determinants of productivity in a group of railroads. The model they use is a two-stage one that, in the first stage, looks at the effect on cost and productivity of characteristics of the chief executive officer (CEO), the effect of the particular method of corporate governance used, and the effect of variables such as energy prices and the state of the business cycle. The key variable capturing firm cost is the ratio of variable costs to revenues. The higher this ratio, the lower is the efficiency of the railroad. The first-stage econometric results fail to show a significant relationship between the background of the CEOs and firm performance measured by the cost ratio. They do indicate, however, that higher education and longer tenure and experience in the firm seem to affect performance negatively. Other performance variables are track utilization and labor productivity, and they do show positive effects from CEOs with backgrounds in operations, law, finance, and marketing. In general, having a rail background appears to have a negative effect on performance.

The same results failed to show strong performance differences among the railroads, although Burlington Northern, Norfolk Southern, and Union Pacific appear to have done best.

The authors turn next to the second stage of their model, which looks for the effect of the cost and productivity variables (the dependent variables from the first stage) on the companies' rates of return. These second-stage estimates also examine the direct effects of variables such as CEO background. The coefficients are conditional upon cost and productivity. The authors find that gross rate of return is positively affected by high labor productivity and by high levels of new investment. It is negatively affected by CEOs who have backgrounds in railroad operations, law, finance, or marketing, and it is positively affected if the CEO was an internal appointment. The importance of

market discipline is shown by the positive relation between gross rate of return and a strong external board, a high debt-equity ratio, and the adoption of shareholder rights plans.

Having estimated both stages of their model, the authors make combined estimates. This tells them directly how variables such as CEO characteristics affect rates of return, combining the direct effect with the indirect effect through cost and productivity. These combined estimates show a mixed effect for those variables that relate to aggregate experience and human capital. Education has a negative effect on all performance measures, as does a CEO who has been promoted internally. In contrast, the effect is positive for chief executives of railroads within a large holding company who have nonrail backgrounds. The authors' measures of corporate governance and control do not show much effect on performance.

In their conclusion the authors note that U.S. railroads have changed dramatically since deregulation. They have tried to diversify, they have smaller plants, they have less labor, and they carry less freight. Almost all of the railroads undertook some form of financial restructuring. Characteristics of rail management obviously matter, but no clear pattern has emerged from this work about the characteristics of an effective railroad CEO. The authors did find, however, that among successful firms, only one relied on an internal CEO.

Baily, Hulten, and Campbell on Plant Productivity

In the third paper Martin Baily, Charles Hulten, and David Campbell examine the behavior of productivity in individual manufacturing plants. They point out that much productivity analysis implicitly assumes that all plants in an industry are identical and respond in the same way to forces that strike the industry. This is in contrast to the literature of industrial organization, where differences among firms are examined in detail. The authors argue that it is important for productivity analysis to recognize the tremendous differences in levels and rates of growth of productivity among plants, even within the same industry.

Empirical analysis in this area can be helpful for policy concerns. For example, their results indicate that plant closings are common even

in successful industries. They are not only a consequence of an industry in decline. Policies that were intended to prevent deindustrialization could inhibit structural change in dynamic industries. Their analysis is also relevant for antitrust policy. They find strong effects on the productivity of a given plant of the productivity of other plants in the same firm. This is consistent with the idea that a well-managed firm can help improve a low-productivity plant and could be a reason to allow takeovers.

The empirical analysis is based upon the Longitudinal Research Database developed by the Center for Economic Studies of the Bureau of the Census. The data are drawn from the five-year Censuses of Manufactures and the Annual Surveys of Manufactures, where individual plants report on their output, inventories, labor input, book value of capital, materials and energy purchased, and other variables. These data allow the estimation of total factor productivity at the plant level, the construction of a distribution of productivity for the plants in each industry, and, with the addition of price deflators, the calculation of plant and industry productivity growth over time.

The authors propose four alternative ways of thinking about the distribution of productivity among plants. The first view is that the observed distribution is simply the result of random shocks or of data errors. In this case plants that were high productivity in one period would not be expected to still be high productivity in a later period. The second view is that plants that have rapid productivity growth in one period will also have rapid growth in the next. This would lead to an increase in the variance of productivity and would suggest that low-productivity plants would quickly have to exit the industry. The third view is that the distribution of plants reflects the age or vintage of the plants. New plants have high productivity and old plants have low productivity and eventually exit the industry. The fourth view is that there are persistent plant differences in productivity. Plants that are high in the distribution of productivity in one period will be high in a later period.

The results of their analysis support a combination of the first and fourth of these alternative views of the productivity distribution. There is plenty of evidence of substantial errors in the data. True productivity in each plant cannot be known exactly. But there is also very strong

evidence of the persistence of productivity. After weighting for size, the authors find, for example, that almost 76 percent of the plants that were in the top productivity quintile in 1972 were in the top two quintiles in 1977. Similar persistence is found in later periods and, strikingly, very high persistence is found even over 10 years.

Baily, Hulten, and Campbell also carry out a decomposition of productivity growth for each of 23 industries and for the aggregate of all of these industries. They find that over five-year periods, there is little impact on productivity growth of the entry and exit of plants from an industry. The bulk of an industry's productivity growth comes from productivity growth within the plants that stay in the industry. There is, however, an important contribution to growth that comes from changes in the shares of output in different plants. The relative importance of this increases when looking at all 23 manufacturing industries together. The shift of output shares toward high-productivity plants and away from low-productivity plants is a major source of productivity growth for manufacturing as a whole.

The authors next develop a regression analysis to examine how the observed characteristics of the plants affect their relative productivity within their industries. They find a strong effect on relative productivity in a given year from relative productivity in a prior period. This reinforces the earlier findings of persistence in the productivity distribution. They also find a strong effect of the productivity of other plants in the same firm. Productivity in a given plant is high if productivity in the other plants of the same firm is high. This could be because of common causes of productivity that operate within a given firm or because of spillover effects—well-managed firms or firms with strong technology passing these benefits on to all their plants. The authors find some tendency for relative productivity to be higher in large plants, but in combination with results from other estimates, they conclude that this does not indicate increasing returns to scale at the plant level. Instead, it appears that large plants may have more market power or a superior product technology, or simply that high-productivity plants may have grown larger.

They find a slight tendency for old plants to be less productive, but the effect is small. They note that age of a plant may be a very weak indicator of the technological vintage of the equipment within the plant.

The authors also examine the relative productivity of entrants and the determinants of the closure of plants. They find in the latter case that size is the largest factor that reduces the probability of closure. Low relative productivity increases the probability of closure.

In their basic regressions the authors find that relative production-worker wages are strongly correlated with relative productivity. High-productivity plants pay high wages. They probe this correlation by asking whether high wages are indicative of high levels of skill or whether workers in high-productivity plants are able to demand high wages. They estimate a two-stage regression in order to separate out these alternatives, and they find that the results favor the latter interpretation. Workers in high-productivity plants do seem to be able to demand higher wages. This interpretation is supported by the fact that new entrants do seem able to use high wages to attract more skilled workers and that high wages do seem to increase the probability of plant closure. Noting limitations of the labor data, the authors suggest caution in interpreting these results.

In their final empirical analysis Baily, Hulten, and Campbell look also at the determinants of the growth of productivity within plants. The most interesting result they find is that a given plant grows more rapidly if other plants in the same firm also grow rapidly. This supports the idea that there are spillover benefits from being part of a successful firm.

Although their data do not have direct evidence on management quality, the authors argue that their results are clearly consistent with the view that high-productivity plants are those that are run by good management teams. Good managers will seek out the best technology, select the most profitable products, and operate their plants efficiently. Their plants will show up as persistently high productivity plants.

Boozer, Krueger, and Wolkon on Race, School Quality, and Wages

Michael Boozer, Alan Krueger, and Shari Wolkon examine differences in school quality for black and white students since the landmark desegregation decision in 1954, *Brown* v. *Board of Education*. Prior evidence has suggested that disparities in school quality between blacks

and whites historically have been responsible for part of the gap in earnings between the races. And since there is recent evidence that the returns to skills have increased in the 1980s, this means that any continuing gap in school quality could provide an explanation of the widening of the racial wage differential that has occurred over the same period.

The authors develop evidence on racial differences in the pupil-teacher ratio, the extent of computer use in the schools, and other school-quality measures since 1954. They concentrate on the resources available to schools rather than on test scores because test scores have been found to be related only weakly to subsequent earnings.

Boozer, Krueger, and Wolkon first address the question of how integrated American schools are, and they find that in 1989–90 the average black student attended a school that was 65 percent nonwhite, while the average white student attended a school that was only 17 percent nonwhite. This pattern of segregation shows up particularly in large cities, where blacks attend highly segregated schools.

Somewhat surprisingly, the authors do not find that this pattern of segregation leads to larger class sizes for black students. The pupil-teacher ratios are roughly the same for blacks and whites. The ratio is higher for Hispanics, largely because this group is concentrated in California where class sizes are generally larger than in the rest of the country.

The computer revolution in the U.S. economy in the 1980s had major effects on American schools. In 1988, 1.5 million microcomputers were available for instructional use in public schools, one computer for every 27 students. Boozer, Krueger, and Wolkon find that black students are much less likely to use computers in school than are white students, even after accounting for family income and other factors. And the gap in computer usage did not decline in the 1980s.

The authors argue that the segregation of schools and the differences in computer use by black and white students may translate into differences in subsequent labor market performance. Black students who attend racially isolated high schools earn less on average than do whites. And students that use computers in schools are paid more than are students without computer skills.

The authors do not believe that school quality is the main explanation for the decline in the relative economic position of blacks in America

since the mid-1970s. The gap between black wages and white wages has widened over this period most dramatically for young workers. And the differences in school quality were smaller for these workers when they were students than was the case for earlier cohorts. Nevertheless, remaining differences in school quality and in access to computers represent a continuing disadvantage for black students when they enter the labor market. Education policy should encourage the development of computer skills among minority students.

Lakonishok, Shleifer, and Vishny on Pension Fund Management

Josef Lakonishok, Andrei Shleifer, and Robert Vishny have obtained access to a proprietary database that has been collected by SEI, a financial services company that advises pension plans in their selection of fund managers. The data are used to evaluate the performance of pension fund managers. Lakonishok, Shleifer, and Vishny use the data to estimate how well these fund managers do relative to the *Standard and Poor's 500 Index*, whether some fund managers are consistently better than others, and how the industry is structured.

Tax-exempt equity funds (mostly pension funds) managed $2.2 trillion at the end of 1990. More than 80 percent of the pension funds are in defined-benefit plans that promise the employee a retirement benefit that does not depend directly on the performance of the fund (provided there is adequate funding). Any residual return not needed to pay retirees is paid to the fund sponsor (the company or organization whose employees are being covered). In defined-contribution plans, by contrast, the amount paid into the fund is set, but the eventual size of the pension depends on the performance of the funds. Since the defined-benefit plans are the more important quantitatively, the authors concentrate on these plans.

Generally, it is in the interest of a corporation with a pension program to ensure that the fund receives the highest return to its assets subject to an acceptable level of risk. Typically, the corporate treasurer of a company will allocate the assets among internal and external fund managers and among different types of funds, such as actively managed and passively managed funds.

Lakonishok, Shleifer, and Vishny argue that the corporate treasurer

will often choose active fund managers as a way of demonstrating that he or she is actively involved in achieving the best results for the company and the employees. And the treasurer will generally choose an external manager in order to reduce responsibility in case of poor performance. Treasurers will even hire outside consultants to select the external managers.

The external managers compete for the business of company pension plans by promising superior performance for a given level of risk. In practice, the fund managers say that they will beat the *S&P 500* by 200 to 400 basis points (100 points equal one percentage point). In principle, the corporate treasurers or their consultants have a lot of information available to them about the performance of the different managers and how they operate, so they can test the promises of the fund managers. The fund managers, however, have some freedom to reveal information that is favorable to them, rather than all information, and to manipulate portfolio choices to enhance their reputations.

The authors turn next to an analysis of their data, which consists of two databases: the performance database and the search database. The performance database, a sample of 769 equity funds, gives quarterly returns from 1983 to 1989. The search database, which is used to evaluate an entire money management organization, not just a particular fund, includes some information back to the 1970s. Since management funds volunteer for inclusion in the search database, the authors judged that there was a significant selection bias involved in it, and they decided to focus primarily on performance since 1983.

The authors compare the performance of the managed funds to the return from the *S&P 500*. They look at the return before management fees are subtracted, and they exclude the cash component of the funds' assets. They conclude that the fund managers have underperformed the *S&P 500* by 1.3 percent a year on average. In fact, weighting the funds by size shows an even bigger degree of underperformance—2.6 percent a year. The fees average about 50 basis points a year, reducing the relative return of the managed funds even more.

Another test of the performance of the managed funds is to ask whether active management pays off. First, the authors compared the performance of the funds, assuming that their assets remained frozen for 6 or 12 months, with the actual performance, given the trades that were made. They found that the trading activities of the funds were

counterproductive. Somewhat surprisingly, however, when they looked at the extent of turnover within funds, they found that higher turnover is associated with higher return. They found that among active managers, those who trade more do better than those who trade less.

Looking at whether there is any consistency of performance, the authors report that funds that have performed well over a two-year period are more likely to perform well over the next two-year period. These findings also hold up over three-year comparisons. There is some consistency of performance and hence some basis for picking funds on the basis of past performance. Some managers appear to be more skillful than others, although even the best money managers cannot be expected to beat the *S&P 500*, net of fees.

Given their weak performance, how do the fund managers survive? One group of fund-managing firms are large banks and insurance companies. Typically, they provide "generic" products such as index funds and annuities. This segment of the industry is concentrated and stable. The second segment of the industry consists of managers who develop a concept and a story. They interact with the corporate treasurers and other company personnel. This segment of the market is much less concentrated, and there is a lot of turnover as treasurers look for the highest fund performance.

In their search for the best performers, corporate treasurers do change fund managers frequently, but there has been only a modest shift away from actively managed funds toward index funds. This is surprising, given the superior performance of the *S&P 500*. The authors argue that treasurers are reluctant to shift to index funds because they then have nothing to do. Treasurers usually diversify the pension fund's assets among several fund managers who have different investment strategies. This gives protection in case of difficulties.

Another possible reason for the continued allegiance of corporate treasurers to active managers is that they interact with these managers and can exercise their own skills at selecting the best managers and the best investment strategies. The fund managers develop differentiated strategies to attract the treasurers, and they develop stories about how their investment strategies are the best. They will offer good excuses to treasurers if performance turns out to be below expectations.

In conclusion, the authors note that the survival of the fund management industry is a puzzle. Here is an industry that seems to be

producing a product with negative value. Eventually, they believe, competitive pressures will force changes in the industry.

Jorgenson, Slesnick, and Wilcoxen on the Effect of a Carbon Tax

In the final paper in this issue, Dale Jorgenson, Daniel Slesnick, and Peter Wilcoxen examine the potential impact of a carbon tax on the economy. Given current concern about the increased concentration of carbon dioxide in the atmosphere, many nations are considering policies to reduce emissions of this gas. Economists generally think a carbon tax, levied on fossil fuels, is the most efficient approach to reducing carbon emissions. In earlier work Jorgenson and Wilcoxen examined the overall effect on the economy of a carbon tax that would stabilize U.S. carbon emissions at 1990 levels. In this study the authors consider not only the overall costs, but also the distributional effects among households. A carbon tax will have its most severe effects on households that devote the largest share of their income to energy products.

Other writers have looked at the distributional effects of a carbon tax, but Jorgenson, Slesnick, and Wilcoxen argue that they have been able to extend prior studies by allowing for full intertemporal general equilibrium effects. They have also developed a method for aggregating the welfare calculations for different households in order to compute a social welfare total. This allows them to examine the impact of a carbon tax on both efficiency and equity.

The model of the economy that they develop divides producers into 35 industrial sectors. The prices of the outputs of each sector depend on the prices of capital, labor, energy, and materials used in production. The price of energy depends upon the prices of coal, petroleum, electricity, and natural gas. The authors have constructed a time series of interindustry transactions tables in order to estimate these price responses econometrically. This approach can be contrasted with "calibration" to data for a single time period used in almost all applied general equilibrium models. Within their production model they generate endogenous productivity growth, so productivity can be affected by changes in the economy, including the carbon tax itself.

The model of consumption involves households that make maxi-

mizing choices over time with rational expectations. In the first stage, households allocate consumption intertemporally. In the second stage, they choose the mix of leisure and goods they wish, and this determines labor supply. The third stage involves the choice among 35 commodity groups, and in this stage the authors allow for differences among household types.

Investment decisions also reflect full maximizing decisions with rational expectations. The price of capital is always consistent with intertemporal equilibrium in asset markets. It is assumed that capital is malleable and that it can be reallocated among the different industries. Capital goods are produced by individual commodities. Based on econometric estimates, the model allows for shifts among types of capital goods in response to price changes, including those that may result from the carbon tax.

To complete the model, the authors must add a government sector and foreign trade. They estimate tax revenues, specify the deficit exogenously (a gradually declining deficit), and obtain total spending, which is then allocated to different commodities according to historical patterns. Imports are assumed to be substitutes for similar domestic goods, while exports are determined by a set of export demand equations.

Carbon dioxide emissions are calculated by taking the carbon content of each fuel type and multiplying by the amounts of each fuel used under different simulations of the model.

The authors turn next to the simulations of the model that will tell them how a carbon tax will affect the economy. They start with a set of default assumptions that determine the exogenous variables in the absence of specific policy changes. The results of this simulation then become the "base-case" results. After that, they impose the carbon tax, where each fuel type is taxed according to its carbon content, and rerun the model. They then compare the results and see how the tax affects the model economy.

Looking first at the long-run effects of the carbon tax, they find that to hold emissions at 1990 levels a tax of $17.65 per ton of carbon is required, corresponding to $11.46 per ton of coal, $2.41 per barrel of oil, and $0.29 per thousand cubic feet of natural gas. This reduces total energy use by about 12 percent and the production of coal by 25 percent.

In the transition to the long run, economic growth has been reduced

as a result of the carbon tax. Productivity growth is reduced by 0.02 percent a year over the period 1990 to 2020, and by 2020, GNP is about one-half of a percentage point lower than it would have been without the tax.

Looking at the distribution of the cost of the tax across households, the authors find that larger households are more affected than are smaller households; households in the West are less affected than those in other regions; farm households are affected less than others; nonwhite households are more affected than white households. All the effects are small, however, averaging about a quarter of a percent of lifetime wealth. This figure rises to 0.325 for nonwhite, low-wealth households and falls to 0.215 percent for unattached individuals with high wealth.

The authors' final calculation computes the effect of the tax on social welfare, allowing for the increased inequality caused by the tax. They estimate that social welfare declines by between 0.15 percent and 0.20 percent, depending on how much weight is put on inequality. Most of the cost of the tax is the result of the loss of efficiency in production; the cost resulting from increased inequality is very small, they conclude.