Summary of the Papers

THE SEVENTH MEETING of the Brookings Panel on Microeconomics was held in Washington in June 1993. The conference papers examined U.S. industrial competitiveness and international comparisons of productivity. Bart van Ark and Dirk Pilat estimated and compared the productivities of the manufacturing sectors of Germany, Japan, and the United States. Martin Neil Baily showed that regulation and the extent of competition affect the productivities of service industries in the United States, Europe, and Japan. Robert Lawrence and Matthew Slaughter made the case that foreign trade is not the reason for the widening of the wage distribution among U.S. workers. Richard Caves and Matthew Krepps examined the hypothesis that companies may suffer from "fat," where they hold too many nonproduction workers. Bronwyn Hall found that the return to research and development (R&D) fell in the 1980s, and she explored the reasons for this. In the final paper Paul Romer presented a new proposal for a voluntary levy on companies to fund technology development.

Van Ark and Pilat on Cross-Country Productivity Differences

The first paper, by Bart van Ark and Dirk Pilat, compares productivity in the U.S., German, and Japanese manufacturing sectors. The authors use new census data and an industry-of-origin approach to show productivity differences and how they have changed over time. The authors are also able to break down the aggregate figures and show how productivity varies by sector. Some parts of the Japanese manufacturing sector are well ahead of the corresponding U.S. industries. The authors also explore the key causes for the productivity differences: capital, scale, and labor skills.

The industry-of-origin approach to productivity comparisons has been developed over the last ten years at the University of Groningen as part of the International Comparisons of Output and Productivity project. As applied to manufacturing, this approach allows estimates of value added per hour worked, measured in national currencies on a consistent basis. The latest year for data in all three countries is 1987.

The most difficult element in any international comparison of productivities is the conversion of the values added in national currencies to a common currency—most studies, including this one, convert all figures to U.S. dollars. Current exchange rates are not suitable for this purpose because most goods, even in manufacturing, are not traded or are subject to trade restrictions of some kind. Instead, van Ark and Pilat use unit value ratios, prices that are calculated using data from the censuses. The unit value ratio is the value of factory shipments of some product divided by the number of units shipped. The ratio of the price of the product in one country to the price of the comparable product in the United States then gives the purchasing power parity (PPP) exchange rate for that product.

The authors were not able to match the entire output of the manufacturing sectors directly on this basis. In many cases reasonable product matches across the countries are impossible to find because the products are unique or differ greatly in quality or specification. The authors find matches covering about one quarter of manufacturing output, and they extrapolate the resulting PPPs to the remainder of production using products from the same industry or branch of the manufacturing sector.

Each of the three countries prepares data on the growth rate of labor productivity in its manufacturing sector. This additional information can then be combined with the 1987 benchmark for relative productivity to give a picture of how productivity has moved over time and the presence or absence of productivity convergence. Because these data are also available by branch (or broad industry level) of the manufacturing sectors, the pattern over time at the branch level can be given.

In contrast to the views of the general public, van Ark and Pilat find that labor productivity in Japan was 78 percent of the U.S. level in 1990; German labor productivity was 86 percent of that in the United States. German productivity was fairly uniformly below the U.S. level across branches, except for basic and fabricated metal products. In Japan productivity in parts of manufacturing (textiles and food, bever-

ages, and tobacco) was only a fraction of that in the United States, while machinery and equipment productivity was well above the U.S. level.

Over time, there has been significant convergence of productivity in Japanese manufacturing to the U.S. level. The productivity gap between German and U.S. manufacturing narrowed through the early 1980s (to the point where German productivity was close to the U.S. level). Since then, however, German manufacturing has experienced a decline of relative productivity.

The authors explore some alternative explanations of the productivity differences. They examine the capital stocks of the three countries and find that the United States had a somewhat higher capital stock per person employed in 1990 than the other two countries. As a result, they find that total factor productivity in Germany and Japan in 1990 was 85–86 percent of the U.S. level. Some part of the explanation of higher U.S. labor productivity is the higher capital intensity of the U.S. sector.

Turning to education, they find that once the importance of vocational training is accounted for, educational differences do not provide an important reason for productivity differences. There are differences in productivity by sector of manufacturing, so that mix effects can account for productivity differences across countries. Accounting for these mix effects lowers German productivity a little and raises Japanese productivity a little relative to the United States. Overall, the authors conclude that capital and education explain only a modest fraction of the differences in labor productivity.

In commenting on the paper, Dale Jorgenson argued that the authors had neglected the role of intermediate goods in productivity, and he presented productivity comparisons from his own work that differed from those in this paper. The authors responded that they had adjusted for intermediate goods in food processing, but that in general the difficulties of estimating PPPs made such adjustments unreliable.

Baily on Competition, Regulation, and Productivity in Service Industries

Martin Baily examines the regulatory and competitive environments in which four service industries operate in the United States, Europe,

and Japan. He argues that these economic environments in Europe and Japan have reduced productivity in the four industries (airlines, banking, telecommunications, and general merchandise retailing).

Economists have long favored policies to sustain competition, but Baily suggests that the benefits of competition and the costs of regulation may have been underestimated. Industry structures change over time, often resulting in large increases in productivity. If regulation or inadequate competitive intensity impedes the evolution of industries, the productivity costs may be very high.

Baily notes that the service sector is much larger than manufacturing and hence contributes heavily to overall productivity. It is important to average living standards to ensure that the economic environment facing service industries is conducive to their successful performance.

The review of the regulatory environments starts with airlines, where Baily points out that the countries of Europe and the United States never agreed on a comprehensive system of regulation for international air travel. The result has been a mixture of bilateral and multilateral agreements rather than an open-skies policy. A typical bilateral agreement specifies the airlines that can fly a particular route—often one airline per country. The most restrictive rate-setting procedure involves double approval, in which either country can reject a proposed rate change. In practice rates are usually set to provide a reasonable return on capital for the airline with higher costs.

Not all airline agreements are so restrictive. For example, much more competition has developed on the North Atlantic route since the United States moved toward deregulation. But most European airlines still operate with heavily protected markets for large parts of their networks.

Baily turns next to banking. The German industry is made up of a multiplicity of banks, many of them very small. This system of banks was created as a matter of policy after World War II with the goal of achieving "socially responsible free markets." The banking authorities saw a need to avoid cutthroat competition, a philosophy that still prevails. Limits have been placed on competition through regional groupings and agreements under which banks specialize in different financial products.

In the United Kingdom the large clearing banks operated in the 1950s under an explicit cooperative agreement. Even after this was ended, U.K. banks continued to avoid competition, rarely seeking new cus-

tomers by offering higher rates on saving accounts or lower rates on loans. Today, the U.K. banks have a natural forum for continued tacit cooperation because they operate the payments system (unlike the situation in the United States, where the payments system is operated by the Federal Reserve).

The telecommunications industries in Europe and Japan have historically been state-owned monopolies (Japan and the United Kingdom have recently privatized their operations). Differences in strategy among the countries, reflecting policy choices, have greatly affected the evolution of the industries. France, after suffering with an inadequate and antiquated system for years, made a big technology push. Germany and the United Kingdom still use obsolete technology. Japan has a technologically advanced system, but telecommunications has traditionally been seen as a source of employment.

In most countries general merchandise retailing is highly regulated. In some countries regulations control entry, while in most countries regulations cover zoning, opening hours, marketing techniques, and the interactions between retailers and suppliers. Regulation is particularly stringent in Japan, where small local retailers are protected against the entry of larger stores or new store formats.

Having explored the aspects of the regulatory or competitive environment that are likely to have resulted in inefficiency, the paper then examines the estimates of productivity prepared by the McKinsey Global Institute for these four industries.

In airlines functional productivities, measured by output per employee, were examined for the various elements that make up an airline's operations, such as flying the planes, airport handling and maintenance, and so on. These data revealed that, on average, the output per employee for European airlines was only 72 percent of the U.S. level in 1989, with lower productivity in all of the functional areas. Several causes of the productivity difference were evaluated, including the effect of hub and spoke operations, but none provided an adequate explanation of the productivity gap. Partly as a residual explanation, but based also on benchmarking studies of individual airlines, it seems that European airlines have excess employment in many of their activities. State ownership and weak competition have led to inefficiency.

Productivity in the banking industry was estimated by separating bank output into three components: transactions, servicing of savings

and deposit accounts, and credit services. This follows the procedures of the U.S. Bureau of Labor Statistics. Output per employee was higher for all three activities in U.S. banks than in U.K. banks, and the same was true for all but the servicing of savings accounts in the comparison between the United States and Germany. Weighting up the three activities revealed that productivity in Germany was 68 percent of the U.S. level in 1989; the figure for the United Kingdom was 64 percent.

Again, several reasons for the differences were examined. The small scale of German branch offices did hurt productivity there. The greater use of electronic funds transfer in Europe gave an advantage to U.K. and German banks, however. Overall, the greater productivity of U.S. banks stemmed primarily from their greater use of information technology and more effective use of labor. The less competitive European banks showed clear signs of inefficiency.

In telecommunications, a major difference between the U.S. industry and the industries in Europe and Japan is that U.S. customers make vastly more calls per access line. Any measure that includes the number of calls as a component of output shows a large U.S. productivity lead. Even in terms of access lines per employee, the U.S. industry has high productivity, although France and Japan are as high or higher in this dimension.

Technology clearly plays an important role in explaining the productivity differences in telecommunications. Germany and the United Kingdom had obsolete switching systems. Incentives to reduce costs are also important. The United States had a single company that dominated service for many years and still has monopoly provision of local service, but private ownership and the nature of regulation did encourage productivity improvements. The other countries either had until recently or still have state-owned monopolies. These industries still carry excess employment.

In general merchandise retailing, value added per employee was the productivity measure used. Productivity levels were similar in the United States and Germany, lower in France and the United Kingdom, and much lower in Japan. Productivities were not very different across countries, however, when similar types of stores were compared—a department store has about the same productivity in New York or Tokyo or London. The biggest source of productivity difference was that more productive retailing formats have had more chance to develop in the

United States and Germany. A Wal-Mart or a Home Depot is vastly more productive than the small retailers they displace. Zoning laws and other restrictions have prevented the full development of these formats, especially in Japan.

Baily concludes that the productivity case studies support the initial argument. State ownership in telecommunications and airlines, restrictions on competition in airlines, regulation and low competitive intensity in banking, and zoning and other restrictions in retailing have hurt productivity in these industries in Europe and Japan. The United States is far from perfect in its competitive environment, but it has generally favored more competition and less state ownership than these other countries. The result has been higher productivity in the service sector.

In commenting on the paper, Robert Gordon argued that the value of comparisons of productivity levels is limited without a discussion of the slowdown in productivity growth.

Lawrence and Slaughter on International Trade and Wages

The real wages of U.S. workers have grown very slowly since 1973; in fact, by some measures they have declined. At the same time the inequality of wages has increased rapidly. The wage gap between more highly educated, skilled, and experienced workers and the relatively less skilled is much greater now than it was in the 1970s.

In the minds of many people, these wage trends have been driven by the United States' position in the world economy. The productivity gap between the United States and other countries has narrowed and the volume of imports has increased substantially. Politicians, ordinary citizens, and some economists all blame foreign trade for the changes in economic performance. Robert Lawrence and Matthew Slaughter argue that this popular perception of the effect of trade on wages is incorrect.

Looking first at the growth of average wages, Lawrence and Slaughter point out that some of the most frequently cited wage series are misleading because they do not include white-collar workers, the self-employed, and fringe benefits. Real compensation per hour in the business sector—a wage series that includes all of these elements—grew by 1.5 percent between 1979 and 1991.

The growth in real compensation, however, was less than the growth of productivity. The authors find no evidence that workers' buying power was reduced by international trade. Indeed, between 1979 and 1991 import prices actually rose more slowly than the prices of U.S. exports. Instead, the authors explain the gap between wage growth and productivity growth on the basis of differences in price movements for the products that workers produce and the products that they consume. On the one hand, the relative price of residential housing rose, and housing makes up a much larger share of consumption than of production. On the other hand, relative prices of capital goods fell, and these are items that workers produce but do not consume.

Turning to the issue of increasing wage inequality, the authors note that the literature has reached a consensus that labor supply factors, such as increased immigration of unskilled workers and the slowing of the entry of educated, skilled workers, do not fully account for U.S. relative wage behavior.

This implies that changes in the demand for labor must have been important. If demand has shifted away from lower-skilled workers toward higher-skilled workers, then the relative wages of higher-skilled workers will have increased as a result. And the trends from manufacturing data support this view. From 1979 to 1989 the ratio of nonproduction to production workers rose from 0.35 to 0.44 (Lawrence and Slaughter argue that separating the manufacturing work force into nonproduction and production workers provides a reasonable proxy for skilled and less skilled). That relative wages of nonproduction workers rose even as the relative number of nonproduction workers employed increased indicates that shifts in labor demand must have been at work.

The second part of the paper assesses whether international trade has driven the shifts in labor demand. To answer this Lawrence and Slaughter look first at the traditional theory of trade to see whether this model is consistent with the argument that trade has caused the shift in labor demand. How could trade have done this? Suppose trade had caused an increase in the demand for U.S. goods that use skilled labor intensively and a decrease in the demand for U.S. goods that use unskilled labor intensively. Then the relative wages of unskilled workers would indeed fall, consistent with the observed pattern in the United States. But two other changes would also have occurred. Employment of unskilled workers would have increased in all manufacturing industries (because

unskilled workers are now cheaper). And the relative price in the United States would have declined for the goods that use unskilled labor intensively (for the same reason).

The authors develop an industry data base to examine both of these propositions, including data on industries at a disaggregated level, and they find no support for either proposition. The relative employment trends have already been noted at the aggregate level—the relative employment of nonproduction workers has increased—and the authors find that this pattern applies in most individual industries. On the price side they find that the relative prices of products that use nonproduction labor intensively generally fell.

Lawrence and Slaughter ask next how technological change would affect these results. Using the rate of growth of total factor productivity by industry as their indicator of technological change, they find that this growth was more rapid in the industries that used nonproduction workers more intensively. This pattern has helped shift overall U.S. manufacturing labor demand toward nonproduction workers. This is not the whole story of the impact of technology, however. To explain why most industries have become more intensive in their use of skilled workers, there must also have been bias in technological change toward the employment of skilled workers.

The authors conclude their study by looking at some qualifications to the trade model. They argue that these cannot reverse their overall conclusion that international trade has not been the source of either the slow growth of wages or rising wage inequality.

Robert Hall said that he found the main argument of the paper convincing. He did suggest, however, that the authors' examination of trends in the employment of skilled labor in tradable goods industries neglects the big positive change in the skill composition of the labor force as a whole.

Caves and Krepps on Squeezing out Excess Manufacturing Employment

Lawrence and Slaughter showed that the basic trend in U.S. manufacturing in the 1980s was an increase in the employment of nonproduction workers relative to the employment of production workers.

Caves and Krepps argue that this increase may be a sign that some companies have accumulated too many nonproduction workers at times and become inefficient. They define this excess employment as the accumulation of corporate "fat" and suggest that fat is particularly likely in large companies.

The authors argue that the fat will be driven out if the companies suddenly face increased competitive pressure. And they suggest that this happened in the United States starting in the mid-1980s, when many white-collar employees found that they had less job security than they had thought. Between August 1989 and August 1990, 65 percent of the increase in total unemployed workers were managers, professionals, and clerical workers. Employment also shifted from large companies toward small business. In 1978, 48.6 percent of all workers were in companies with fewer than one hundred people; this had risen to 51 percent by 1984. The number of workers in companies with more than one thousand workers fell from 18.6 percent to 16.2 percent.

A standard assumption of economic theory is that companies maximize their profits, which should ensure that excess employees are not retained. Caves and Krepps argue that in certain situations, firms might employ excess white-collar employees and suffer low white-collar productivity. For one thing, the output of each individual worker is hard to measure, making it difficult for managers to compare wage and marginal product. For another, the objectives of managers may include increasing their own satisfaction by hiring extra staff assistants or other employees.

Caves and Krepps move next to the empirical phase of their study. Controlling for the dependence of employment upon wages and output, they test the hypothesis that competitive pressures can drive out excess employment. They look specifically at whether changes in corporate control and increased import competition can be expected to have negative effects on white-collar employment.

To test the effects of a change in corporate control, the authors build a data base giving the proportion of assets in each industry that were subject to changes in control between 1965 and 1986 and relate that to white-collar employment in the industry. They hypothesize that takeovers in any given industry will reduce the fat in that industry as a whole, possibly through a demonstration effect for companies that are not themselves subject to takeovers.

Caves and Krepps report that for 1967–86, the measures of change in corporate control and the measures of import competition are not terribly significant as determinants of white-collar employment, although the directions of the estimated effects are correct. The authors argue that these initial results are not terribly surprising. Changes in competitive pressure occurred at different times, and thus variations over time must be studied to get a clearer picture. After allowing for changes over time, they find that the effect of mergers was to reduce white-collar employment in the 1980s, although this result did not hold in earlier years. Similarly, the rise in import penetration seems to have had its main effect on excess employment in the 1980s.

The authors look next at differences among sectors. For example, acquisitions of technology-intensive companies might not reduce white-collar employment if the workers were involved in R&D or related activities. On the whole, the results do not indicate a major difference in effects by sector.

The authors judge that their results provide some support for the view that increased pressure forced companies to reduce their white-collar employment. To examine the hypothesis further, they look at data on the stock market's reaction to news of individual corporate downsizings. The announcement of a corporate downsizing will have a negative effect on corporate valuation if it reveals information about a company that the shareholders had not known about previously (the "bad news" effect). It will have a positive effect on corporate valuation if it means that managers are shedding excess labor that they had previously been unwilling to discharge (the "bite-the-bullet" effect).

To investigate these possibilities, the authors collected a sample of announcements of corporate downsizings appearing in the *Wall Street Journal* between 1987 and 1991. The authors found that 60 percent of the 513 announcements they studied led to negative changes in stock prices (mean -5.1 percent), and the remainder were positive (mean 3.7 percent).

Using regression analysis, Caves and Krepps looked for attributes of the downsizing announcements that can separate the bad news effects from the bite-the-bullet effects. For example, an announcement that coincided with a poor earnings report had a negative effect on market valuation. An announcement of charges associated with the downsizing also had a negative effect. Overall, Caves and Krepps found that, on

average, layoffs involving nonproduction workers were seen as creating value for shareholders. This effect is statistically significant, but they note that the explanatory power of their models is not very high.

The authors also examined those firms that had high levels of selling, general, and administrative expenses (overhead costs) per employee. They found that a positive effect on shareholder value for small or moderate downsizings in companies with high overhead costs.

In sum, the authors conclude that their results are consistent with the view that fat, or excess white-collar employment, tends to accumulate in successful companies and that it takes a competitive shock of some kind to remove it.

In her comments on the paper, Michelle White noted that the category of nonproduction workers that the authors use to measure excess employment is actually much broader than the concept of white-collar bureaucrats used in theoretical models of excess bureaucracy. Nonproduction workers include engineers, R&D personnel, and salesworkers, categories of workers not generally included in the theoretical models.

Hall on the Rate of Return of Research and Development

When a U.S. company performs R&D, it is accumulating an asset, a stock of knowledge that is a form of intangible capital. This intangible capital should show up in two ways that can be observed—both the stock market value and the productivity of the company should be higher as the asset is used in production.

Bronwyn Hall starts her paper with two observations about the effect of R&D on U.S. companies that have emerged from her recent empirical studies. First, the value that the stock market appears to place on the R&D that companies perform fell sharply in the 1980s, expecially measured against the valuation the market places on tangible capital assets. Second, the contribution of R&D to productivity in those same firms also declined.

The author notes that one obvious explanation for the declining stock market value of R&D capital is a decline in the supply price of R&D funds. If the R&D tax credit or other tax changes had made R&D funds much cheaper to companies, for example, or if the rate at which the stock market discounts funds had fallen, or both, then companies would

do more R&D. This would drive down the return from the marginal R&D project and lower the estimated market valuation of R&D capital.

Hall is able to resolve quickly that this explanation cannot be the main one, which leaves two other possibilities. Either the marginal product of R&D may have fallen, or the rate of depreciation or obsolescence of R&D capital may have increased. Either new knowledge is not helping companies' bottom lines as much as it used to during the first year that it is acquired, or else the market value of that knowledge is fading away more quickly than it used to.

Before turning to her empirical exploration of these issues, Hall examines recent trends in R&D. She reports that the R&D series prepared by the National Science Foundation shows that the ratio of R&D to sales in U.S. manufacturing rose from 2.2 percent in the 1970s to 3.4 percent in the 1980s and has stayed at about this level. By contrast, Compustat data indicate that this same ratio has been rising steadily from about 1.0 percent in the 1970s to 4.5 percent in 1990. Hall suggests various reasons why this discrepancy has arisen.

Turning to the econometric analysis of the company data, Hall reports that in 1980 a dollar of physical capital (equipment and structures) was valued by the stock market at about $0.60 and a dollar of R&D capital was valued at $1.50. By 1990 the valuation of a dollar of physical capital had risen to $1.20, while a dollar of R&D capital had fallen to $0.40. The author notes that although these results are subject to considerable uncertainty, the pattern is so strong that they are very unlikely to result from pure chance.

Hall posits that in 1980, the owners of companies were signaling that manufacturing had excess capacity (or the wrong kind of capacity). The wave of restructuring corrected this situation, and manufacturing capacity is more in line with companies' needs. Because the market also was placing a high valuation on R&D capital, companies naturally responded by raising R&D budgets. Some fall in valuation was to be expected as a result.

These explanations make sense but do not tell the full story. There was apparently some overshooting with respect to physical capital, because its valuation went so high. More important, the fall in value of R&D capital to only forty cents on the dollar requires further explanation.

Hall uncovers a clue to what has happened by looking at six tech-

nology sectors, where she finds that the valuation of R&D differs substantially by sector. The valuation of R&D in chemicals, autos, metals, machinery, and other industries rose from around zero in the mid-1970s to about parity with other assets by the end of the period. Pharmaceutical R&D maintained a high market valuation, while the value of R&D in the electrical, instruments, computing, and electronics sectors fell sharply in the early 1980s and by 1990 was hardly worth anything on the market.

The fall in overall valuation of R&D therefore seems to have been driven by the fall in these hard-hit sectors. Is this from initial marginal product or from obsolescence? Hall tackles this question by decomposing the stock of R&D capital into two parts, new R&D and the R&D carried over from previous years. She finds that both contributed to the decline in valuation. By 1990 both old and new R&D are valued at only one quarter of their value in the 1970s, with the declines concentrated in the same industries noted earlier. The increased rate of obsolescence of R&D capital seems plausible, given what is known of the increased entry and competition among high-tech industries and the resulting rapid rate of imitation for new products and processes.

Hall pushes her empirical strategy further by separating out the extent to which R&D capital contributes to reductions in costs and increases in revenues—and charts how these contributions changed in the 1980s. The results vary by sector, and the differences are substantial depending on the size of the companies. One important clue to the pattern of market value emerged. The productivity of R&D fell significantly in large electrical and especially computing companies. This decline is associated with the declining fortunes of the large mainframe manufacturers.

Hall notes in conclusion that these patterns are consistent with a widening gap between the private and social rates of return to R&D. Even though new technology may contribute to overall productivity growth, the ability of individual companies to appropriate this return has fallen sharply.

In commenting on the paper, Edwin Mansfield said that in his work with companies, the connection between the amount spent on R&D and the profit obtained was very weak. Success in implementing new product or process technology may be more important, he said.

Paul Romer on a National Technology Strategy

In the final paper, Paul Romer proposes a mechanism to stimulate technological change. He argues that increases in the stock of knowledge have driven economic growth in the past and that the potential for further increases in knowledge is vast. The opportunities for sustained economic growth through enhanced knowledge creation are remarkable, he asserts.

The limiting factor in exploiting technological opportunities is the difficulty of collective action, because it takes collective action to encourage discovery. Romer sets forth a specific institutional arrangement he has developed that is designed to stimulate collective technology development while avoiding the pitfalls of typical government interventions.

Consider a situation where the manufacturers of a particular product agree that research could improve the product or process technology in the industry. The companies in the industry would then apply for government permission to hold an election to impose on all companies in the industry a levy on sales (say, 1 percent). If the proposal passed, then it would become law and the tax would be assessed.

The revenue from the tax would go to industry-organized boards that would fund research—for the sake of example, assume one for university research and one for research on process equipment. Each industry board would be controlled by directors nominated by the industry members. Individual companies would be free to allocate their funds to either board, provided the company's total contribution matched its tax rate. Knowledge created by the research would have to be made available to all industry participants. Foreign companies would be free to participate, based upon their sales in the United States. And the details of how to administer the levy (such as defining the boundaries of the industry) would be determined by the secretary of commerce as the head of the supervising government agency.

Romer recognizes that this proposal will strike some readers as unrealistic, but he points out that it has precedents. The Agricultural Marketing Agreement Act of 1937 outlines procedures for establishing arrangements called marketing orders, which set up marketing boards. About three-quarters of these boards collect funds for research and

marketing. Some boards impose quantity restrictions, which, Romer says, should clearly not be allowed under his proposal. Other similar precedents are provided by the Electric Power Research Institute, which finances research, although not very effectively; and Bell Labs, which was financed by the Bell operating companies before the breakup of AT&T and is judged to have been very successful.

To justify the establishment of a scheme like his, Romer points out that he must show both that there are unexploited gains to be had from collective action and that these are important in magnitude. To explore the first issue he distinguishes between rival and nonrival goods. Ordinary goods are rival because once such a good has been consumed by one person, it is no longer available to others. Knowledge is nonrival because one person can use some piece of information, such as a computer program, without depriving another person of the opportunity to use the same knowledge. It is in the production of nonrival goods that collective action offers the the greatest potential gains.

The usual results showing the efficiency of markets break down for nonrival goods. Once a discovery has been made, efficiency suggests that it should be made freely available—because it is not used up in consumption, there is no need to ration the good. But without restrictions on the use of knowledge, many or all users of the invention will be "free riders," paying nothing for its development. Without restrictions on use, a creator will not be rewarded for making a discovery—the returns from an invention must be appropriable in order to reward the inventor.

In practice many goods are a mixture of rival and nonrival characteristics, and the ability of producers to appropriate returns varies. This means that policy must address the trade-off between the incentives for discovery and those for diffusion or use of knowledge, with the nature of the trade-off varying for particular cases. How is the cost of discovery to be shared and the problem of free riders avoided?

Markets and governments resolve the trade-off in different ways. Government has unique power to solve the cost-sharing problem, but it also has unique power to waste resources. Markets solve the cost-sharing problem only by introducing monopoly, but they are good at selecting opportunities and avoiding waste. For such things as movies or the design of new microprocessors, the market works much better than government. For basic research, the system of government-pro-

vided funds allocated with peer review works pretty well. There is an intermediate zone of generic research, however, where opportunities may be missed by the current system, and it is here that Romer sees the need for new institutional arrangements of the type he is describing. In his proposal, government power solves the free-rider problem, while market participants decide where the opportunities are greatest.

Romer says that according to the evidence, the likely magnitude of the social rate of return to production of nonrival goods—invention or the creation of new knowledge—is on the order of 30 to 50 percent. These are investments with very high payoffs. And this evidence typically ignores innovations that are made without traditional R&D, such as the development of discount retailing.

Romer turns next to the political issues posed by his proposal. He argues that science funding has become subject to serious political interference in Congress and that bureaucracies must overcome inertia and inefficiency to solve problems. One of the advantages that Romer sees for his proposal is that it could change the economic context in which current technology decisions are made. For example, it could change the nature of the interactions among researchers in companies and in universities.

A potential complaint about his proposal, says Romer, is that it may not be suitable for all industries and hence might end up being a small program. This may actually be an advantage, says Romer, given the potential for adverse effects of policies. After all, little is lost if industries simply decide not to participate in the program. There would be no tax levy and no program.

Romer concludes by noting that even after admitting all of the limitations of markets in the area of knowledge creation, the problems of any of the political alternatives must still be reckoned with. He sees his proposal as a way to solve the free rider problem while maintaining the advantages of market-based systems.

In the discussion of Romer's proposal at the meeting, participants expressed concern about the practicality of the plan, given the difficulties of defining industries and hence who would pay into a levy. The potential for collusion created by sanctioning joint industry activities was also discussed.