

Pay, Participation, and Productivity

Alan S. Blinder

America's dismal productivity performance during the 1970s and 1980s is the nation's chief economic disorder and should be the chief concern of economic policymakers. From 1973 to 1988, output per worker-hour in all U.S. businesses grew at a paltry compound rate of 1.05 percent a year. That is barely more than a third of the growth rate we enjoyed during the halcyon 1947-73 period (2.96 percent a year) and, more important, only about half the long-term historic average. Had U.S. productivity grown at 2 percent since 1973, instead of 1.05 percent, standards of living in the United States would now be about 16 percent higher. If the rate of productivity growth stays so depressed for a protracted period of time, America is destined to slip into the second rank of nations in terms of wealth and income, just as the United Kingdom did before us. To most Americans, that is a distasteful prospect.

Suggestions for boosting U.S. productivity are not hard to come by; but *good* suggestions are rare. The best way to raise productivity growth, and perhaps the only way to do so permanently, is to speed up the pace of technological innovation. Unfortunately, no one has a reliable formula for making a society more innovative, although better scientific education and more spending on research and development would presumably help.

By contrast, we know much more — though certainly not as much as we would like — about how to increase the rate of capital formation. Perhaps that is

why both economists and policymakers who think about ways to spur productivity keep coming back to incentives for saving and investment. For example, suggestions for various and sundry tax breaks for capital abound.

Unfortunately, help from this quarter is likely to be small, and for a simple reason: Capital inputs amount to no more than 30 percent of the economy's total costs. A permanent 10 percent increase in the rate of investment, which would eventually raise the nation's capital stock by 10 percent, would be a signal achievement. But with capital accounting for only about 30 percent of total costs, a 10 percent increase in the capital stock would raise labor productivity (output per hour of work) only about 3 percent. And even this small effect would be a long time coming. After five years, a 10 percent higher investment rate would be expected to boost labor productivity only 1.2 percent, about two-tenths of a percent a year.

Labor presents a more tempting target because it accounts for at least 70 percent of total costs. If we could figure out a way to make labor 10 percent more efficient, so that workers could accomplish in an hour what now takes 66 minutes, output per hour of work would rise by about 7 percent even with no increase in capital. Furthermore, such an increase in labor productivity would soon pull investment along. And if that were to happen in a short time, the transitory increase in productivity growth would be impressive.

A nation has many ways to try to improve the ef-

fectiveness of its labor inputs. One of them is changing the way labor is compensated for its efforts. If a change in the pay system could indeed raise productivity, it would be a particularly attractive way of doing so, for it would require little or no sacrifice of real resources. By contrast, other (possibly good) ideas for improving the quality of the work force — reducing illiteracy, for example — carry steep price tags.

Yet people who worry about productivity rarely think or write about changing the way workers are paid. It is as if existing institutional arrangements, including the wage system, were part of the natural order of things. In fact, however, a society starting over again to design a pay system to encourage high productivity would be most unlikely to choose the conventional wage system. Workers are now paid not for output produced, nor even for labor input provided, but simply for time spent on the job. From what economists and psychologists know about incentives, such a system should not be expected to call forth labor's best efforts. The notion that linking pay to performance might improve labor productivity is hardly startling.

What is "obvious," however, is not always true, and introspection is a notoriously unreliable guide to empirical magnitudes. Do incentives like piece rates, profit sharing, and gain sharing have important effects on labor productivity? To begin to answer those questions, Brookings asked me to assemble a group of experts on issues of pay and productivity to examine the empirical evidence available both here and abroad. Their findings, which are summarized here, confirm, more strongly than I expected, that pay incentives can improve productivity. More surprising, the experts find that the way workers are treated may affect productivity as much as, if not more than, the way workers are paid.

Lessons of History

The idea that individual or group pay incentives might motivate workers to higher levels of performance is not new. Piece rates date back to the preindustrial era and were in common use in manufacturing during the 19th century. Both gain sharing, which ties wages to some measure of group output or costs, and profit sharing trace their roots back to the 19th century.

Daniel Mitchell, David Lewin, and Edward Lawler find two noteworthy themes in the history of thinking on and experience with alternative compensation systems. One is that rising and falling tides of interest in the various incentive plans have more to do with changing social, political, and economic fashions than with accumulating scientific evidence on how well the plans work. The other is that advocates of incentives have typically thought of and designed them not as partial substitutes for straight wages, but as add-ons

that raise total compensation. The authors dub this the "gravy" view.

Mitchell and his coauthors offer several pieces of econometric evidence to show that recent incentive plans have been gravy. First, they estimate from industry wage surveys covering the years 1979–86 that workers on incentive plans like piece rates or production bonuses earned about 11 percent more an hour than other workers. Second, using data from the 1970s, they estimate that workers who received cash bonuses (including cash profit sharing) were not paid lower straight wages. Third, using data from the 1950s, they estimate that the presence of profit sharing in the compensation package did not reduce other fringe benefits.

Taken together, these findings invite the interpretation that workers on profit sharing or incentives are more productive than workers on straight wages. After all, if they did not produce more, why would their employers pay them more? But even if workers on profit sharing are more productive, two critical questions remain.

The first one is, do profit sharing and other incentives actually boost productivity, or do they simply attract the most productive workers to jobs where high productivity is rewarded? This issue is critical to em-

The Experts

Alan S. Blinder is the Gordon S. Rentschler Memorial Professor of Economics and chairman of the department of economics at Princeton. The experts whose work he has summarized here are Michael A. Conte, professor of economics at the University of Baltimore; Masanori Hashimoto, professor of economics at Ohio State; Douglas L. Kruse, assistant professor of industrial relations and human resources at Rutgers.

Edward E. Lawler III, professor of management and organization and the director of the Center for Effective Organization in the School of Business at the University of Southern California; David I. Levine, assistant professor, Haas School of Business Administration at the University of California, Berkeley; David Lewin, professor of business at Columbia.

Daniel J. B. Mitchell, the director of the Institute of Industrial Relations at the University of California, Los Angeles; Jan Svejnar, professor of economics at the University of Pittsburgh; Laura D'Andrea Tyson, professor of economics at the University of California, Berkeley; and Martin L. Weitzman, professor of economics at Harvard.

A complete report on their findings can be found in the new Brookings book, *Paying for Productivity: A Look at the Evidence*, which Blinder edited.

pirical research on alternative pay schemes. But it may not be critical to a firm deciding whether to institute profit sharing. After all, such a firm cares more about *whether* its productivity will rise than why it rises. From society's point of view, however, the source of the productivity gain is crucial. If profit sharing merely shifts workers from one company to another, society as a whole neither gains nor loses. But if profit sharing actually raises the productivity of individual workers, society reaps an important benefit.

The second question is whether profit sharing raises productivity enough to pay for itself. At first blush, the answer seems obvious: If worker productivity goes up and profits are shared, owners of capital must come out ahead. But the issue is more subtle under the gray view. For example, suppose workers on profit sharing earn, on average, 10 percent more than workers on straight wages. If a firm is to benefit from profit sharing in that case, labor productivity must rise at least 10 percent. If it goes up less than 10 percent, the firm loses money despite the gain in productivity. Not enough is known about productivity gains from profit sharing, but it is far from clear that profit sharing pays for itself.

Mitchell, Lewin, and Lawler survey a wide variety of case study evidence, largely of rather low quality, and find that most of it supports the notion that profit sharing, gain sharing, and incentive payments help boost productivity. The case histories also suggest that profit sharing and gain sharing seem to succeed more often when they are combined with some type of worker participation in decisionmaking.

The authors also estimate the effects of both alternative compensation methods and worker participation on various measures of firm performance, using the new Columbia Business Unit data set. They find that employee participation, profit sharing for production workers, and possibly employee stock ownership plans (ESOPs) have positive effects on productivity. The estimated average productivity gain from profit sharing is 8.4 percent, which is sizable. But the authors' regressions fail to detect the positive interaction between participation and profit sharing suggested by the case studies.

Profit Sharing and Productivity

Martin Weitzman and Douglas Kruse focus on whether profit sharing increases productivity. To answer that question, they begin with the unremarkable proposition that paying for performance should elicit better performance than paying just for time. Thus profit sharing should be superior to straight wages. This simple intuition seems compelling at first, but it is open to three main theoretical objections.

The first might be called the one of many, or $1/n$, problem. An individual worker in a firm with many employees gains little if his efforts raise company

profits. Realizing this, he or she may be inclined to shirk. If everyone shirks, group productivity is low and everyone fares poorly. If everyone works hard, creating more profits to share, the whole group is better off. But how can the cooperative equilibrium be sustained when n is large?

Weitzman and Kruse note that this problem is a classic prisoner's dilemma game, in which the prisoner knows that the police do not have enough evidence to convict him unless his partner, who is also under arrest, implicates him. But the prisoner also knows that, in the event of a conviction, he can get a lighter sentence by ratting on his partner. The dilemma is whether to keep silent (the cooperative solution) or talk to the police. In the work place, the prisoner's dilemma game is one that workers play repeatedly, not just once. Since repeated games admit a multiplicity of solutions, possibly including the cooperative one, Weitzman and Kruse conclude that the $1/n$ problem, though real, is not a decisive objection to profit sharing. But will the cooperative, high-productivity solution actually emerge? The theory is mum on this point. As the authors put it, managers may have to do more than install profit sharing and walk out the door. To garner a productivity bonus, they may have to develop a corporate culture that emphasizes cohesiveness and cooperation.

The second theoretical problem is risk aversion. Linking pay tightly to performance might lead to higher productivity. But it will also give workers more variable income, which they may dislike. Efficient labor contracts must balance these two concerns, and the theoretical literature suggests that the optimal contract almost certainly mixes some profit sharing with a base wage.

The third issue pertains to the appropriate roles of labor and capital in decisionmaking. Some critics argue that profit sharing leads to demands for worker participation in management and that such participation necessarily undermines efficiency. In theory, however, efficiency demands that capitalists make all the decisions only if they can supervise and monitor workers costlessly, a situation most unlikely in practice. If, on the other hand, workers can monitor and motivate fellow workers more effectively than managers can, then worker participation may actually raise productivity. Weitzman and Kruse suggest that theory reinforces the finding that the optimal labor contract probably has a profit-sharing component.

Economic theory thus modifies and qualifies, but does not overturn, the simple intuition that tying pay to performance should raise productivity. What does the empirical evidence say?

Weitzman and Kruse examine a broad array of evidence. First, a look across both socialist and capitalist countries suggests to them that productivity is probably higher when more people care about making en-

terprises profitable. Second, like Mitchell, Lewin, and Lawler, they find that the case study evidence shows that profit sharing has a positive effect on productivity and that worker participation may help. Third, opinion polls usually find that both employees and employers believe profit sharing boosts productivity.

Fourth, and most intriguing, Weitzman and Kruse note that a large number of independent, though relatively weak, results can add up to a strong statistical case when the results point in the same direction. A literature search turned up 16 studies using 42 different data samples that estimated the productivity effect of profit sharing. Many of these studies have flaws; none are beyond reproach; several obtain weak results. But the consistency of the disparate results is striking.

Out of the 218 estimated profit-sharing coefficients, only 6 percent are negative, and none significantly so. By contrast, 60 percent of all the regression coefficients are significantly positive. The odds that this could happen by chance are infinitesimal. This, I believe, is the strongest evidence adduced to date that profit sharing boosts productivity.

How large is the estimated productivity dividend? Among the 101 coefficients for which Weitzman and Kruse have enough information to answer the ques-

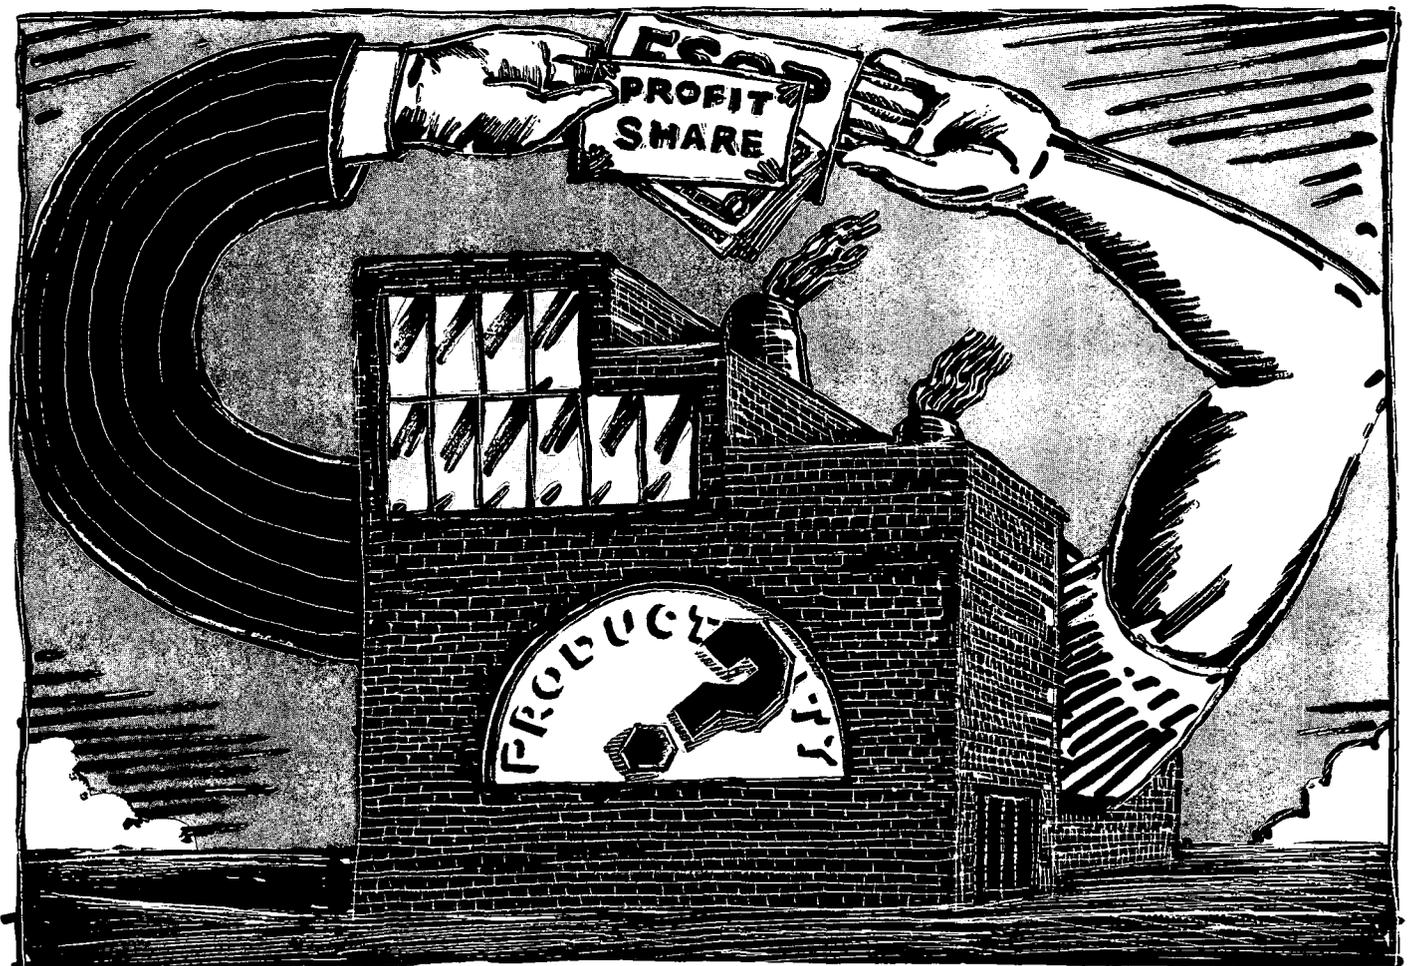
tion, the median estimate is 4.4 percent, and half of them fall between 2.5 percent and 11 percent.

Employee Stock Ownership Plans

Employee stock ownership plans are sometimes confused with profit sharing. In principle, they are something quite different. ESOPs pay benefits in company stock rather than in cash, and the company's contributions need not be tied to profits. Deferred profit sharing plans, however, are more like ESOPs, especially when the deferred compensation is held in company stock. And most profit sharing is deferred.

ESOPs were accorded special tax status largely as a way to democratize American capitalism by spreading stock ownership, not as a way to improve worker performance. Nonetheless, the argument that tying the fortunes of labor and capital together might improve productivity has been applied to ESOPs, just as it has to profit sharing. In fact, many of the theoretical arguments dealt with by Weitzman and Kruse apply equally well to ESOPs.

After reviewing the limited number of studies on ESOPs and productivity, plus some related studies of European cooperatives, Michael Conte and Jan Svejnar judge the evidence too mixed to support any



strong conclusion. ESOPs have not been found to harm either productivity or any other measure of company performance, and several studies suggest that they help. But the strong positive effect estimated in the first study of ESOPs and productivity has not been verified by subsequent research. On the other hand, Conte and Svejnar argue that the evidence is much stronger that participatory institutions combined with ESOPs improve performance. They believe it is not voting rights or seats on the board of directors, that help the most, but in-plant committees and consultative mechanisms.

Appraising Worker Participation

Literature surveys suggest that worker participation boosts the effectiveness of pay incentives. But do participatory institutions foster higher productivity?

Solid empirical evidence on this question is hard to come by because participation is a vague term encompassing a wide variety of labor-relations practices, including quality circles, work teams, labor-management consultation, and employee representation on corporate boards. Moreover, the existing studies use divergent methodologies, and many offer no quantitative measure of the effect of participation on productivity. Nonetheless, David Levine and Laura D'Andrea Tyson make a heroic effort to summarize the disparate empirical evidence.

If one ignores research on cooperatives — which are special (though often successful) cases — there are 29 studies in all. Of these, only 2 conclude that participation hurts productivity, whereas 14 find that it helps notably or in a lasting way or both. The other 13 offer more ambiguous results; the estimated productivity gains are small, short-lived, or dependent on other factors. This evidence makes it much easier to believe that participation helps rather than hurts productivity, Levine and Tyson conclude.

Under what conditions do participatory arrangements work best? Levine and Tyson argue that there are four mutually supporting pillars:

- *profit sharing or gain sharing*, to make workers feel that cooperative behavior is rewarded;
- *guaranteed long-term employment*, to give workers long time horizons so they do not feel threatened by change;
- *relatively narrow wage differentials*, apart from those attributable to seniority, to promote group cohesiveness and solidarity; and
- *guarantees of worker rights*, such as dismissal only for just cause.

If participation really can raise productivity at little cost to firms, why is it so rare in the United States? Why is it so much more common in some other countries? Levine and Tyson offer two possible answers.

The first is that the attractiveness of participatory institutions depends on the economic environment

within which firms operate. Participation works better in countries where average unemployment is low, recessions are mild and infrequent, and providers of capital have long time horizons and close working relationships with the businesses they finance. (According to Levine and Tyson, the causation works both ways: participation also helps achieve low unemployment and mild recessions.) These conditions characterize Japan and Sweden, say, much better than the United States.

The second explanation rests on externalities: participatory arrangements work best for one firm when they are also used by others. For example, a firm that pays narrower wage differentials than others may find its best workers quitting, while less-experienced, less productive applicants flood its personnel office. The problem evaporates, however, if all firms offer the same egalitarian wage structure. Similarly, a firm that dismisses only for just cause may find itself attracting an inordinate number of workers who like to shirk, if competing firms continue to dismiss workers at will. But if all firms adopt dismissal only for just cause, the adverse selection problem disappears. The upshot, Levine and Tyson argue, is that an economy may have two equilibriums: a superior one with participation, and an inferior one without.

Worse yet, if attained, the participatory equilibrium, may have a tendency to unravel. Consider wage differentials, for example. If, starting from an equilibrium with narrow wage differentials, one firm begins to offer premium wages to star workers, other firms will find themselves losing their best talent. If they react by widening their own pay differentials, the participatory equilibrium comes apart. A strong social consensus or legislative encouragement may be necessary to achieve and maintain an equilibrium with high levels of participation.

Labor Relations in Japan

A few years ago, many Americans thought that the key to Japan's industrial success was more and better capital. Now we are coming to realize that the Japanese edge comes from knowing how to use people, not machines. That is certainly Masanori Hashimoto's view. He points to nine significant differences between the Japanese and American industrial relations systems:

□ Long-term employment is both more prevalent and more formal in Japan, where dismissals and layoffs are extremely rare. So the Japanese worker has enough job security to identify with the company and to accept industrial change willingly.

□ Wages are tied more tightly to seniority in Japan than in the United States, which Hashimoto interprets as evidence that Japanese companies invest more in firm-specific human capital.

□ Japanese workers receive a substantial share of

their compensation in semiannual bonuses, which are uncommon in the United States. These bonuses make wages more flexible and, to the extent that they serve as profit sharing, might enhance productivity.

□ Japanese unions are organized along enterprise lines, not along the craft or industrial lines that typify American unions. The Japanese model creates greater commonality of interests between labor and business leaders and makes it easier for unions to assist in monitoring worker performance.

□ Wage bargaining in Japan is synchronized in an annual spring offensive called *shunto*. This system helps make wages more flexible than in the United States, where collective bargaining is decentralized and staggered.

□ Japanese labor contracts are brief and sketchy, leaving a great deal to on-the-spot decisions in the workplace arrived at by mutual consent. American labor contracts, by contrast, are long documents that try to spell out everything in writing, thereby rigidifying the employment relationship.

□ Quality control circles are used much more, and apparently more effectively, in Japan.

□ Japanese managers devote much more time and energy to joint consultations with labor than their American counterparts do. These procedures, which are often formalized, facilitate communication, foster harmony, and promote a fair sharing of the gains from innovation.

□ Decisionmaking within Japanese firms is based on consensus building, including the involvement of labor. American managers exercise more control from the top.

Hashimoto's list goes well beyond the four factors emphasized by Levine and Tyson. Which are the key ingredients? Hashimoto contends that joint consultation and consensus-based decisionmaking are the cornerstone of the Japanese industrial relations system. Japanese businesses make enormous investments of time and money in these mechanisms; they must believe that their investments pay off. Why, then, do American businesses not do the same? There are only three possibilities: American business is underinvesting in participatory mechanisms, or Japanese business is overinvesting, or such investments are more profitable in Japan than in the United States.

Hashimoto favors the last explanation. He is therefore hesitant to recommend importing Japanese industrial relations practices to America. Specifically, he hypothesizes that the Japanese invest more in the employment relation because the costs of joint consultation and of reaching decisions by consensus are lower in Japan, whereas the benefits are more or less the same in the two countries.

Is the hypothesis true? We don't know. We do know that Japanese managers take more care in screening job applicants. But that just shows that they invest more in the employment relation, not that it is cheaper

to do so. Hashimoto suggests that the greater homogeneity of the Japanese labor force reduces the costs of communication, which is plausible. But the fact that Japanese management techniques seem to work well in the United States — with American workers — gives one pause.

Levine and Tyson emphasize the first explanation, which leads them to recommend more participation. But the second explanation — that the Japanese overinvest in employee relations — should not be ruled out. Upper and middle managers in large Japanese companies spend many hours after work with employees in restaurants, at company picnics, and so on. Often the company picks up the bill. If these hours were counted as part of the workweek, rather than as leisure time, Japanese labor productivity, as conventionally measured, would be lower than it is.

Thus, though the differences between Japanese and American industrial relations practices are clear, the reasons for them are not. Nor is it clear that American firms would do better by moving toward Japanese industrial relations practices — although many people, including myself, suspect that this is true.

Conclusions

What, then, has this look at the evidence told us. First, there is good reason to believe profit sharing does indeed raise productivity, but much less reason to believe ESOPs do so. (Gain sharing may be the best system of all, but there are too few cases to support any strong judgments.) The evidence on profit sharing and productivity is particularly persuasive if one accepts the view that a large number of weak, but consistent, studies add up to a strong statistical case. It is not known, however, whether the productivity dividend is typically large enough to pay for the profit sharing, if the profit sharing is given as gravy.

Second, worker participation apparently does help make alternative compensation plans like profit sharing, gain sharing, and ESOPs work better — and also has beneficial effects of its own. This theme was totally unexpected when I organized the conference.

Which forms of participation raise productivity the most? We do not know. But giving labor a seat on the board of directors may be the least effective form of employee participation. Beyond that, there is controversy — and insufficient statistical evidence to resolve the dispute.

So, it appears that changing the way workers are treated may boost productivity more than changing the way they are paid, although profit sharing or employee stock ownership combined with worker participation may be the best system of all. That finding, and the evidence behind it, should be of intense interest to policymakers concerned with raising productivity and to business executives concerned with raising profits.

Stemming the Spread of Chemical Weapons

Elisa D. Harris

"Assessing the proliferation of chemical and biological weapons is one of the most difficult challenges we face in the intelligence community now and into the next decade. It is also one of our most important tasks, for these weapons may well represent one of the most serious threats to world peace in the coming years."

— CIA Director William H. Webster, October 25, 1988

Concern over the spread of chemical weapons has mounted steadily since the mid-1980s when Iraq used poison gas against Iranian military forces and later against its own Kurdish population. The potential dangers inherent in the proliferation of these weapons became even more apparent following revelations of Libya's attempts to acquire the capability to produce chemical weapons.

Elisa D. Harris, a senior research analyst in the Brookings Foreign Policy Studies program, is currently working on a study on chemical and biological arms control. This article is based on a paper Harris prepared as an SSRC-MacArthur Foundation Fellow in International Peace and Security Studies for a conference sponsored by the Aspen Strategy Group in Aspen, Colorado, August 13–18, 1989. A fully documented version will appear in the forthcoming conference volume, New Threats: Responding to the Proliferation of Nuclear, Chemical and Delivery Capabilities in the Third World, to be published by the University Press of America and Aspen Strategy Group.

These incidents sparked several actions during the past year to strengthen and expand the international legal regime governing chemical weapons. In January 1989, 149 countries reaffirmed their support for the 1925 Geneva Protocol prohibiting the use of chemical weapons in war. In September, the United States and the Soviet Union announced initiatives aimed at enhancing the prospects of a global treaty for the actual elimination of chemical weapons. First, at a meeting in Jackson Hole, Wyoming, the two sides agreed to exchange information and carry out inspections of each other's chemical weapons stockpiles and production facilities. And then at the United Nations, President George Bush proposed that the superpowers begin unilateral but mutual reductions in their chemical weapons stockpiles.

A halt to the spread of chemical weapons cannot come too soon. Already 12 countries outside NATO and the Warsaw Pact are thought likely to possess or to be developing chemical weapons. Other countries have been reported to be working toward acquiring such a capability. Who are these countries? What are the implications of continued proliferation? And what is the best potential solution to the problem?

The Data Problem

Despite numerous press reports alleging the development, possession, or use of chemical weapons by various countries, detailed and reliable information sub-