

the right as well as to the left in the United States, and the advocates of ESOPs have exploited this appeal quite successfully.

Moreover, lawyers, accountants, investment bankers, and management consultants—the actors in society principally responsible for the design of business organizations—have long organized themselves in employee-owned firms. They cannot be unaware of the benefits of employee ownership or opposed to it on principle. At most they can be accused, rather implausibly, of hoarding the benefits of employee ownership for themselves and—whether out of spite or just lack of imagination—denying those benefits to firms in other industries.

To understand the prevailing pattern of employee ownership, we must turn instead to the costs of governance that are the subject of the next chapter.

6

Governing Employee-Owned Firms

A recurrent theme in the voluminous literature advocating employee ownership (or, more broadly, “economic democracy,” “worker participation,” or “labor management”) is that employee participation in control of the firm through democratic processes is of value in itself, quite apart from the quality of the substantive decisions reached by those processes. In Chapter 3 we speculated on three reasons why this might be so: that participation in governance is a consumption good; that it provides a valued sense of control; and that it stimulates and informs participation in political life beyond the boundaries of the firm. The last of these reasons might provide some justification for public subsidy to employee ownership as a means of making workers more responsible citizens in a democratic society. Unfortunately, the available empirical evidence provides little support for it.¹ The employees themselves enjoy the other two potential benefits of employee participation. If those benefits are actually important, they should influence employees’ choices about the types of firms in which to work. That is, they should give employee-owned firms a survivorship advantage.

Participation is not free, however. It brings with it all the costs of collective decision making. And there is substantial evidence that these costs can be large.

The Costs of Collective Decision Making

In many respects a firm’s employees are often better situated than its investors to oversee management effectively, as we observed in Chapter 5. But there is a compensating disadvantage: employees are far

more likely than investors to differ among themselves concerning the firm's policies.

To begin with, employees may disagree about their relative wages. Further, employees often have different stakes in the firm's investment decisions, such as which plants to keep open, which processes to automate, or where to improve safety. The extent to which employees' interests diverge in these respects is likely to increase as the division of labor within the firm increases. Where all employees do essentially the same job, there is little reason for disagreements about pay; moreover, the employees will be similarly affected by most decisions and thus generally in agreement about them.

Employees can also experience conflicts of interest that have other sources besides differences in job assignments. For example, employees can differ substantially in the amount of equity capital they have invested in the firm. This is particularly likely to occur if, as is common, the firm's pension fund is the principal vehicle for employee investments. Older employees, who have disproportionately large amounts of capital invested, will prefer to have a larger amount of the firm's earnings attributed to capital (and hence distributed as earnings on amounts invested in the pension fund) and a smaller amount attributed to labor (and hence paid out as wages) than will younger employees.

Such conflicts might not be troublesome in practice if there were obvious objective criteria to employ when making the decisions in question. For example, if the actual marginal product of both labor and capital within a given firm could be easily measured, then it would be natural to apportion the firm's earnings between wages and return on capital in proportion to those marginal products. Such objective criteria, however, are usually absent or unobservable at reasonable cost. In most firms, for instance, the marginal productivity of labor and capital cannot be accurately measured. Therefore, in important decisions, there is often considerable room for judgment and discretion, and hence for active disagreement. And disagreement, as discussed in Chapter 3, can make governance costly.

We observed in Chapter 4 that, while different groups of shareholders in investor-owned firms sometimes have conflicting interests, corporate law provides means for constraining and resolving the worst of these conflicts in corporate decision making. It is reasonable to ask whether the law could place analogous constraints on collective decisions made by employees. One place where this has been tried is under

the "duty of fair representation" that American labor law imposes on unions to prevent the exploitation of the minority by the majority. But in contrast to the analogous doctrines in corporate law, this body of law has been singularly unsuccessful in limiting opportunistic behavior by majorities. In general, the duty of fair representation has been effectively employed only to bar overt discrimination based on conventional personal characteristics such as race or sex.² The reason for this is apparently not that the law has had inadequate opportunity to develop, but rather that, unlike the situation with investors, there are no simple objective criteria by which to determine whether a particular subgroup of employees is being treated unfairly.

When we survey the types of firms in which employee ownership has succeeded, the costs associated with collective decision making appear surprisingly important in determining where employee ownership is viable and how it is organized. In fact, these costs seem to go far toward explaining the large residual in the existing pattern of employee ownership that is left unexplained by the other considerations surveyed in Chapter 5.

Which Firms Succeed?

The most striking evidence of the high costs of collective decision making is the scarcity of employee-owned firms in which there are substantial differences among the employees who participate in ownership. Most typically, employee-owners all do extremely similar work and are of essentially equivalent status within the firm. Rarely do they have substantially different types or levels of skills, and rarely is there much hierarchical authority among them. This is evident in the professional service firms, where employee ownership is best established. The partners in a law firm, for example, are all lawyers of roughly equal skill and productivity who work more or less independently of each other; rarely does one partner have substantial supervisory authority over another. Similarly, the employees in the U.S. plywood cooperatives are unspecialized and only semi-skilled, and commonly rotate over time through the various jobs in the mill. They thus have little reason to differ concerning the policies to be adopted by the firm. The manager, who is the only person in the firm with specialized skills and tasks and with substantial supervisory responsibilities, is generally not a member of the cooperative but rather hired as a salaried employee.³

The driver-owned transportation cooperatives that are so common throughout the world also fit this mold; the drivers have extremely similar stakes in any decision made by the firm.

The predominance of this pattern suggests strongly that employee ownership works best when there is little opportunity for conflict of interest among the employee-owners. Evidently the viability of employee ownership is severely compromised when the employees who share ownership play diverse roles within the firm⁴ and consequently will be differently affected by important decisions taken by the firm—a conclusion supported by substantial anecdotal evidence from a variety of different types of employee-owned enterprise.⁵

Conversely, employee ownership evidently is viable when the employees involved all play a similar role within the firm, even if in other respects the net benefits from employee ownership do not appear unusually strong in comparison to situations in which investor ownership is the rule. As we saw in the last chapter, employee ownership does not typically arise where the costs of contracting for labor on the market are unusually high; indeed, quite the contrary. Apparently employee ownership is worthwhile, as a means of reducing the costs of labor contracting, even where those costs are, compared with those of other industries or other types of employees, relatively low—as long as collective decision making is not a potential problem. This remains true, moreover, where the risk-bearing costs associated with employee ownership are relatively high.

The natural conclusion from this pattern is that, if costs associated with collective self-governance were not a problem, employee ownership would be far more widespread than it is.

Structures to Avoid the Costs

Another important indication that collective governance can be costly, and that the costs involved play a critical role in the success or failure of employee ownership, lies in the strong tendency for employee-owned firms to adopt rules and practices that promote homogeneity of interest among the employee-owners.

For example, the plywood cooperatives nearly all adhere rigidly to a scheme under which all members of the firm receive the same rate of pay regardless of task or seniority. The firms explicitly justify this practice as necessary to avoid excessive dissension among the members.⁶

Even more striking, many of America's largest and most prosperous law firms have long followed a practice of sharing the partnership's earnings equally among all partners of a given age, regardless of individual productivity. This is an astonishing fact. As we have seen, law firms not only can, but do, monitor the productivity of their individual lawyers quite closely. It would therefore seem natural to adjust each partner's return to reflect her productivity, and thus provide a strong financial incentive for efficient performance. Yet the equal-sharing firms abandon all such financial incentives. Obviously there must be some strong countervailing value served by equality in distribution of earnings.

One explanation that has been offered for this practice is that it is a mechanism for risk sharing.⁷ Yet it is hard to believe that risk sharing is the principal motivation. At the time they enter into the equal-sharing scheme, the lawyers involved have already achieved great prosperity and security. They have proved their professional competence and have become partners, with effective lifetime tenure, in an established law firm. Commonly their expected earnings are at least several hundreds of thousands of dollars per year. Could such individuals nevertheless be so risk averse that they are prepared to relieve their partners of all financial incentives for productivity simply to ensure that their own income will always be the same as that of the other partners their age?

This is highly implausible.⁸ It is more likely that these equal-sharing schemes serve principally to reduce the costs of collective decision making. An equal-sharing rule provides a simple focal point for deciding how to divide the pie. Using a political process to decide on a more complex differentiated scheme of division would be time-consuming and divisive for all involved, and there is no reason to believe that a stable outcome could be easily achieved.

Law firms that do not adopt equal-sharing rules commonly employ formulas under which a partner's share is determined according to specified indicia of productivity, such as hours billed or number and value of new clients brought to the firm. Such formulas—as opposed to less formal approaches under which a manager or committee is simply given discretion to set relative shares as seems appropriate, without being bound to a rigid formula—are evidently an alternative effort to establish more or less objective, and hence uncontroversial, criteria for dividing the pie where equal sharing is too difficult to justify. Even so,

there is considerable dissension within law firms about the structure of these formulas, and the resulting disagreements are an important source of instability and dissolution among law partnerships.

Indeed, employee ownership generally seems to thrive only where, if equal sharing is not practicable, individual employee productivities are sufficiently easy to measure so that some relatively objective, and hence uncontroversial, method of pay based on that measure can be employed. We find employee cooperatives among taxi drivers and refuse collection crews, where members of the cooperative bill clients individually and can simply be compensated with a fraction of those billings. Employee ownership is correspondingly rare among those firms—which constitute the overwhelming bulk of all large firms today—in which production requires the joint effort of large numbers of employees with different skills performing different tasks, so that the productivity of individual employees is difficult to assess with any precision.

Employee-owned firms also commonly strive to ensure that not only pay, but also amount and even type of work, is equalized among the members of the firm. The employee-owners in the plywood factories, as already noted, commonly rotate through the different jobs over time, so that there is little long-run specialization of work among them. Law firms strongly resist admitting to the partnership any lawyer who is not of roughly the same competence and productivity as the other partners; less-qualified lawyers, if valuable to the firm, are kept on as permanent salaried associates rather than as partners who receive a smaller share of earnings.⁹

Similarly, law firms strongly resist letting some partners work fewer hours than average in exchange for a smaller share. The recent rapid increase in the number of women lawyers, for example, has created considerable pressure for part-time work arrangements to permit time for child rearing. Many law firms now willingly accept such arrangements for salaried “associates” (young lawyers not yet promoted to partnership), and at the same time refuse to permit women to be partners on a part-time basis.¹⁰ This refusal is sometimes explained on the ground that clients demand that attorneys be available full time or that attorneys must practice full time to keep up their skills.¹¹ But these explanations seem forced. Rather, it appears likely that such inequalities among members of the firm are resisted at least in part because they tend to destabilize the cooperative governance structure. A simple

rule under which everyone does essentially the same amount and kind of work, and receives the same pay, is by far the easiest to agree upon and to enforce,¹² and these advantages evidently are often sufficient to outweigh the costs such a simple rule engenders in the form of inflexibility, poor incentives, and lack of diversification among the work force. (This is not to deny, of course, that simple sexism also plays a role in denying partnerships to women under conditions acceptable to them.)

To be sure, it is possible that such egalitarian practices are unusually common in employee-owned firms at least in part for other reasons. In particular, it could be that causation runs the other way from that suggested here: employee-owned firms may be unusually inclined to adopt egalitarian practices simply because they are employee-owned, and not because such practices are necessary to reduce governance costs. For example, there is evidence that employees judge the adequacy of their salaries in relation to the amount paid other employees at the same firm. Consequently, an employee will accept a lower wage if he is among the best paid at the firm, and will require a larger wage if he is among the lowest paid. The result is that the wage structure within a given firm is less dispersed than differences in productivity among employees would predict; the employees at the top of the wage scale compensate the employees at the bottom, as it were, for the privilege of being on top.¹³ Perhaps this effect is intensified when employees are co-owners rather than simply salaried employees, because they are then even more inclined to view themselves as a collective reference group for purposes of judging their individual welfare. If so, this would lead employee-owned firms to have less differentiated wage structures than similar investor-owned firms.¹⁴ Such a phenomenon would not, however, explain why it is that employee ownership tends to arise only where the employees involved are highly homogenous to begin with.

In recent years, the size of corporate law firms has increased dramatically. Firms are now highly departmentalized, and the size of teams that work on major projects has also become large. As a result, there is now substantial horizontal and vertical division of labor within firms. Under such circumstances, norms requiring equal contributions from all senior attorneys become harder to maintain, and equal sharing of earnings becomes harder to justify. These developments presumably help explain the increasing tendency among law firms to hire “perma-

ment associates” who remain with the firm indefinitely as salaried employees rather than as partners. By this device, ownership of the firm can be confined to a relatively homogeneous class of attorneys without cramping the firm’s growth or diversification. For similar reasons, it seems predictable that law firms might seek to accommodate as permanent associates rather than as partners the increasing numbers of senior women attorneys who wish to work only part time, although laws and norms concerning sex discrimination may inhibit firms from pursuing this approach.

The ultimate means for a service firm to avoid the governance costs associated with employee ownership, however, is to convert to investor ownership. And this is, in fact, the path being followed in many of the service professions. Advertising firms began converting from the partnership form to investor ownership in the early 1960s, and most of the larger firms are now investor-owned.¹⁵ Similarly, investment banking began abandoning the partnership form in the 1970s, and many of the larger firms there, too, are now investor-owned.¹⁶ Medical practice has been following the same route, as investor ownership has spread rapidly among health maintenance organizations.¹⁷

Evidently one reason such firms convert to investor ownership is to attract larger amounts of capital than can conveniently be supplied by the professionals who work for them. But there is evidence, too, that the conversions serve to alleviate governance problems. Firms in these industries have been growing larger and more complex, offering a broader range of services and exhibiting more internal specialization and departmentalization. Presumably consensus among the professionals within these firms concerning policies to be followed and division of the firm’s earnings has become increasingly difficult to secure. Striking evidence of this appears in the well-documented sale of Lehman Brothers, one of America’s oldest investment banking partnerships, to American Express in 1984. Although a need for capital seems to have played a role, the sale was precipitated by a breakdown in internal governance that had its roots in conflicts between the old-line bankers and the newly powerful traders within the firm.¹⁸

Further evidence comes from the recent emergence of a number of “boutique” investment banking firms organized as partnerships.¹⁹ These firms have only a small number of partners and concentrate on only a portion of the investment banking business (such as mergers and acquisitions), hence reducing opportunities for serious conflicts in de-

cision making. At the same time they are often exceptionally well capitalized.²⁰ This indicates that the need for capital is not a decisive factor in the choice of ownership in investment banking, and suggests strongly that governance also plays an important role.

What Constitutes Homogeneity of Interest?

The preceding evidence implies that employee ownership works best where the employee-owners are so homogeneous that any decision made by the firm will affect them roughly equally, or where, though the employees differ in ways that cause the burdens and benefits of some decisions to be shared unequally, there is an objective and widely accepted basis for making those decisions. That is, employee ownership is most viable where either no important conflicts of interest exist among the employee-owners, or some simple and uncontroversial means is available to resolve the conflicts that are present.

What constitutes a difference, and what constitutes an objective and legitimate way to resolve a difference, are of course relative things. Formal public choice theory confirms what common sense suggests: that strong homogeneity of interest is not in itself sufficient to eliminate the pathologies of collective choice mechanisms. If the range of available outcomes to choose from is sufficiently large, then majority rule may not yield a stable outcome even if preferences are only very slightly heterogeneous.²¹ To take a mundane example, if a majority of the pilots in a pilot-owned airline have brown eyes, they might vote for a policy under which brown-eyed pilots get paid 20 percent more than do blue-eyed pilots. But then, if the pilots over five feet ten inches in height form a majority, they might subsequently vote to eliminate the eye-color-based pay policy in favor of a policy in which pilots over five feet ten inches get paid 30 percent more than do shorter pilots. And so forth.

We must therefore ask why voting ever works. Why doesn’t shared decision making among a group of owners always break down, in favor of either chaos or authoritarianism, when there exists any basis for distinguishing among the owners, no matter how trivial—as there virtually always is among a firm’s employees? If, as the evidence suggests, substantial (but of course never complete) homogeneity of interest is sufficient to reduce considerably the costs of collective decision making, there must be some special reason why this is so.

One answer, presumably, is that culture and institutions place important limits on the range of legitimate voting outcomes. Where there is substantial perceived homogeneity of interest, as where all workers have essentially the same role within the firm, socially shared norms of fairness, as well as the fiduciary rules that a reviewing court will enforce, place all but a few potential outcomes beyond reach. Thus neither conventional morality nor the courts are likely to sanction a system in which pilots with brown eyes get paid more than do those with blue eyes.

It follows that homogeneity of interest is, in important degree, a social construct. Consequently, what passes for homogeneity in one setting may not in another. There was a long period in which it would have been both morally and legally acceptable in the United States for pilots with white skin to be paid more than pilots with black skin. As we shall see in Chapter 7, an important reason why farmer-owned marketing cooperatives were for a long time much less common in tobacco than in other crops was apparently the presence of much more racial diversity among tobacco farmers than among other types of farmers, with the result that shared governance was less workable among the tobacco farmers.

Ownership by a diverse constituency can be made more viable by creating and maintaining institutions—including shared norms, formal legal and charter limitations on voting procedures and voting outcomes, and review by outside agencies such as courts—that constrain the available range of choices.²² We have already seen a number of examples. The costs of constructing and sustaining these institutions, which increase as the perceived degree of heterogeneity among the owners increases, are an important component of the costs of collective decision making.

Representative Democracy in Practice: The Case of Mondragon

The firms in which employee ownership is found are, like law firms, often small enough to permit the use of highly participatory forms of direct democracy. This is not to say that such firms are, in reality, usually highly participatory. In partnerships of professionals, such as law firms, although the partners may have substantial nominal rights to participate in collective decision making, governance is often effec-

tively confined to a small group of senior partners who are essentially self-appointing—a practice that itself suggests that collective decision making among employees is costly.

In any event, in larger employee-owned firms, highly participatory forms of decision making are unwieldy, and direct democracy must usually yield to a representative form of governance. The most obvious model is one, analogous to that employed in widely held investor-owned corporations and in other types of large cooperatives, under which the employee-owners elect representatives to a board of directors. The board, in turn, is responsible for appointing and overseeing the firm's managers. Direct voting by the owners is then confined to major "constitutional" changes in the firm, such as merger or liquidation. Where the work force is heterogeneous, indirect representation of this sort could have an advantage over more direct forms of democracy in securing professional management and avoiding high process costs or inefficiently biased decisions. Yet such a system of indirect representation might also insulate management so well as to sacrifice some of the potential advantages of employee control.

Few well-documented examples of large employee-owned firms with heterogeneous work forces exist, and so it is difficult to assess directly and in detail how well a representative system of employee governance can function in such a setting. The case that has been most extensively studied is the affiliated group of worker cooperatives at Mondragon, Spain. Because of the substantial success these firms have achieved (by a variety of measures their average productivity has regularly exceeded that of Spanish industry in general),²³ and because of the extraordinary amount of attention they have received among advocates of employee ownership, they call for special consideration. Unfortunately, even the literature on Mondragon does not focus very carefully on matters of governance. We cannot draw strong conclusions, consequently, but rather can only point to some aspects of the Mondragon experience that offer evidence of the costs of decision making.

In an individual firm within the Mondragon group, direct employee participation in governance is largely confined to annual meetings at which the employees elect representatives, in at-large elections, to a nine-member supervisory board of directors for staggered terms of four years. The board, in turn, is responsible for appointing the firm's managers, who serve for a minimum of four years and cannot be removed during that term except for cause.²⁴ (In contrast, the mem-

bers of the board of directors of American business corporations are typically elected for terms of only one year, and managers serve at the discretion of the board.) Each firm also has a "social council" that is elected by the employees separately from the elections to the board of directors and that serves as the principal avenue by which the interests of the employees in a firm are made known to management. In contrast to the members of the board of directors, who are elected at large, the members of the social council are elected by local constituencies within the firm. The social council, however, acts only in an advisory capacity to management and has no formal authority.²⁵

In addition to these procedural rules, which blunt the use of electoral mechanisms as a direct means of controlling management, there are limitations on the substantive decisions that the employees can make. In particular, the employees in an individual firm within the Mondragon group are constrained in the extent to which, and the ways in which, they can appropriate the firm's net earnings. They are not free to set their wages at any level they wish; rather, each firm must adopt a system of wages that deviates only within specified bounds from a scale established by the Mondragon system's central bank.²⁶ Moreover, 10 percent of the net earnings remaining after payment of these wages must be devoted to educational, cultural, or charitable purposes;²⁷ the rest must be retained and invested in capital accounts. At least 20 percent of net earnings must be put in a collective account that cannot be appropriated by the firm's employees, even upon retirement from the firm, and this percentage increases according to a formula as the firm becomes more profitable. The remaining profits—no more than 70 percent of total profits—are invested in accounts for the individual employees. The amounts in these accounts cannot be withdrawn before retirement but earn interest (which is paid to the workers annually in cash) at 6 percent. These financial restrictions are imposed on the individual cooperatives by a "contract of association" that they enter into, at the time of their formation, with the Mondragon system's central bank; the restrictions are apparently not subject to alteration by a cooperative's employees or management.²⁸ Other potentially contentious topics, such as precedence in hiring and layoff, are also governed by systemwide rules.²⁹

This attenuation of employees' rights to earnings and control in the individual cooperatives is reinforced by other features of the system. Individual cooperatives that are engaged in a given type of production

—consumer goods manufacturing, for example, or furniture making, or agricultural production—are federated into a larger organization whose leadership is effectively appointed by the managers of the constituent cooperatives. These higher-level organizations often have considerable authority over their member firms. In particular, there is substantial pooling of profits among the firms within a given group, so that profitable firms underwrite the losses of unprofitable ones. Similarly, individual workers can be reassigned from one firm to another within the group.

In addition, each individual cooperative must affiliate itself with the system's central bank. The bank, whose own board of directors is largely composed of managers of the constituent cooperatives, exercises substantial authority over the individual firms. The bank's wage-setting authority has already been mentioned. Each firm must also obtain whatever additional capital it needs from the bank and must invest any capital surplus with the bank.³⁰ The bank also retains and exercises the authority to replace an individual cooperative's management, or to take over the firm's operations directly, in case of poor performance.³¹

This system, with its strong delegation to management and its centrally imposed restrictions on individual firms' decision making, has evidently been designed self-consciously to mute the play of political forces among the firm's employees, preempting opportunities for costly conflict. Not surprisingly, given this structure, the Mondragon firms have a somewhat managerial character. It appears that leadership comes largely from the top; for the most part, the electoral mechanisms are employed largely to ratify proposals made by management.³² The powerful bank at the center of the system itself appears distinctly managerial, as late as 1987, the two individuals who had long served as chairman of the bank's board of directors and as its chief executive officer were both among the five persons who had founded the Mondragon group more than thirty years before, and the chairman of the bank's board was also chairman of the board of the oldest and largest individual firm in the Mondragon system and of the federated group to which that firm belongs.³³ An eight-day strike of workers against managers at the latter firm, in 1974, is indicative of the managers' independence from their workers.³⁴

In short, workers' rights to control and to participation in earnings are attenuated in the Mondragon system, and the individual firms

cannot really be said to be fully owned by their workers. Rather, the system has something of the character of a large nonprofit holding company that delegates to the employees of each subsidiary the right to elect the subsidiary's management. More particularly, Mondragon bears some resemblance to the more decentralized American private universities, such as Harvard, whose constituent schools and departments have substantial budgetary autonomy and are largely self-governing but are subject to the right of the central administration to intervene when it feels necessary.

None of this is to deny that the employees at Mondragon participate meaningfully in control. An indication of that control (though perhaps also of some instability in governance) is the departure from the group of at least four of the cooperatives through conversion to investor ownership, and the departure of at least another twelve firms as independent cooperatives. All of these departures took place despite the opposition of the central administration.³⁵

The wage structure has been among the most contentious issues at Mondragon. Initially, the systemwide top pay index—the maximum permissible ratio of highest to lowest rate of pay—was three to one. Over the years, this ratio has been increased to attract and hold talented executives. Thus in 1988 the top pay index was raised from 4.5 to 6. The new top index was imposed only as a ceiling, however; individual cooperatives were left free to adopt a lower ratio if they wished. The change was unpopular with the rank and file, which repeatedly voted down acceptance of the higher top index at individual firms. Five years later, in 1993, only 8 of 125 affiliated firms had adopted the new top index of 6, and nearly three-quarters remained at 4.5 or below. (Even the latter firms appear liberal as cooperatives go: three-quarters of Spain's other—generally much smaller—producer cooperatives report that, as with the American plywood cooperatives, all their members receive the same wage.)³⁶

In sum, Mondragon does not offer a straightforward demonstration of the feasibility of adopting full employee ownership in a large firm with a heterogeneous labor force. But it does show that an attenuated form of employee ownership, in which the firm is managed on behalf of its employees though not fully owned by them, can sometimes be successful in industrial enterprise.

The representativeness of the Mondragon experience, as many have noted, remains unclear. There are reasons to be cautious in concluding

that replication will often be feasible. Among these reasons are the ethnic homogeneity, insularity, and low mobility of the Basque population from which the Mondragon system draws its work force.³⁷ After forty years in operation, and despite substantial publicity, Mondragon has still spawned few successful imitators on a similar scale, either in Spain or elsewhere.³⁸

In any event, examination of Mondragon reinforces the inference that the costs of collective decision making are significant in employee-owned firms, and that structures to constrain those costs are a strategic element in the design of those firms.

Other European Experience

Mondragon aside, the characteristics of worker cooperatives throughout Europe—for all one can tell from the literature, which tends to be scanty on significant details—seem to follow a consistent pattern that reinforces the observations made above concerning the efficiency of employee ownership.

Among European countries, as we have noted, Italy and France apparently have the highest concentrations of successful worker cooperatives.³⁹ In the case of Italy, this may be at least in part because that country (like Spain) subsidizes worker cooperatives in general through tax exemptions and special credits.⁴⁰ In addition, both France and Italy grant construction cooperatives special advantages in bidding for government business, which may be quite important in explaining the high concentration of cooperatives in that industry.⁴¹ The construction companies and artisanal firms that constitute the bulk of the worker cooperatives in these countries seem to conform roughly to the type of firm in which employee ownership has proved most viable elsewhere, with relatively little internal hierarchy, although the Italian cooperatives evidently include some construction companies and manufacturing firms of substantial internal complexity.⁴²

The Italian worker cooperatives are generally affiliated with one of three associations, each of which has traditionally been connected with one of the nation's political parties. In particular, well over half of the cooperatives, including most of the largest and most successful firms, are members of an association (the "Lega") historically affiliated with the former Italian Communist Party (now the Democratic Party of the Left). Although it is difficult to obtain a clear view of these coopera-

tives from the available literature, there is some reason to believe that their decision-making autonomy is constrained to an important degree by control exercised from the top down through the Lega.⁴³ Moreover, the cooperatives' members are represented by a national labor union also affiliated with the former Communist Party, and wages in the cooperatives, like those in similar investor-owned firms, are determined by industrywide bargaining—with the Lega serving the role of the employers' representative for the cooperatives in these collective negotiations. Furthermore, in Italy (as in France), the statutes governing worker cooperatives impose on them a quasi-nonprofit structure, requiring that a substantial fraction of net earnings be retained in accounts not distributable to the worker-members and prohibiting the distribution of net assets to members upon dissolution.⁴⁴ Profit-sharing bonuses are limited by law to 20 percent of wages. In practice, these bonuses are commonly paid to nonmember employees as well as to members, are often small or nonexistent, and (like wages) do not exceed those paid in investor-owned firms.⁴⁵ In sum, although the evidence suggests an important degree of real employee governance in the most successful Italian worker cooperatives, those cooperatives operate in an environment of favoritism and constraints that make it difficult to draw general conclusions from them about the behavior and viability of fully worker-owned firms.

England has far fewer employee-owned firms than France or Italy. Omitting partnerships of service professionals, such as solicitors, the firms that have been successful over the long run can be divided into two groups. The first group has its origins in the labor and cooperative movements of the late nineteenth century. The number of these firms has been declining throughout this century; by one count there were only sixteen remaining in 1973. The firms in this group are typically small artisanal cooperatives; more than half are in printing or bootmaking.⁴⁶ The second group consists of more recently founded firms associated with the Industrial Common Ownership Movement (ICOM). There were eleven of these firms as of 1977. They include some moderate-sized industrial enterprises with a fair record of success. The ICOM firms, however, are not truly employee-owned. The employees have no claim on the firms' net assets, their rights to participate in current net earnings are constrained, and ultimate authority over important aspects of the firms' affairs are in the hands of trustees who are not elected by the employees.⁴⁷ These firms are, like many other firms

called worker cooperatives, essentially nonprofit firms that are managed on behalf of their workers and not actually owned by them. They have also benefited from some generous subsidies. Four of the eleven ICOM cooperatives, for example, were essentially donated to their employees by their former owners.⁴⁸

The experience in these countries thus does not provide clear evidence of the viability of large firms with heterogeneous work forces under full employee ownership.

Experience with Partial Employee Participation

The importance of governance costs relative to the other costs and benefits of employee ownership is further underscored by the accumulated experience with various forms of partial employee participation in earnings and control that fall short of full employee ownership.⁴⁹

Employee Stock Ownership Plans

Since the 1970s many American firms have adopted employee stock ownership plans under which most or all of the firm's employees receive a portion of their compensation in the form of stock in the firm. Typically ESOPs are structured as deferred compensation plans in which the employer deposits stock in a trust fund that holds the stock for the benefit of the participating employees, often as the reserve for the employees' pensions. By 1986 approximately 4,700 companies had adopted such plans. Roughly 25 percent of these plans owned more than 25 percent of the stock in their firms, and somewhat less than 2 percent owned all of the stock.⁵⁰ By 1990 the number of companies with ESOPs had grown to roughly 10,000,⁵¹ and estimates suggest that the ESOP owned a majority of the stock in more than 1,000 of these companies.⁵²

The rapid proliferation of ESOPs is not an unbiased indicator of their efficiency. Although the ESOP concept has been actively promoted since the 1950s,⁵³ it did not become popular until ESOPs were granted substantial federal tax subsidies beginning in 1974—tax subsidies that have since been broadened and deepened⁵⁴—and until it was discovered that creation of an ESOP could be a useful defensive tactic for management in an attempted corporate takeover.⁵⁵ It is entirely

possible that, without these special advantages, ESOPs would remain rare. The numerous studies of ESOPs that have been undertaken to date, while not conclusive, have failed to present clear evidence of improvements in either employee productivity or firm profitability once allowance is made for tax subsidies.⁵⁶

Whatever the motivation for adopting ESOPs, one of the most striking facts about them is that they generally provide for participation only in earnings and not in control. Only rarely have they been structured to give the employees a significant voice in the governance of the firm. A substantial fraction of the stock held by ESOPs is non-voting stock.⁵⁷ Further, the power to vote any voting stock held by an ESOP is commonly not exercised by the employees who are the beneficiaries of the plan. In the latter regard, the tax law plays a significant role. For a privately held corporation to obtain the tax benefits provided for ESOPs, the power to vote stock held by the corporation's ESOP need not be passed through to the employees, but can be voted by the plan's trustee.⁵⁸ The trustee, in turn, can be appointed by the firm's management without consultation with the employees who are the plan's beneficiaries. In publicly traded corporations, in contrast, voting power must be passed through to the employees on all ESOP stock actually allocated to the employees—which is to say, not purchased through borrowing, as in the popular "leveraged" ESOP.⁵⁹

These tax law provisions are evidently important in understanding the pattern of ESOPs that has evolved. If we exclude so-called tax credit ESOPs—that is, ESOPs created under a special (and now repealed) provision effectively giving a 100 percent tax subsidy to the plan—roughly 90 percent of all ESOPs are in privately held firms. Moreover, there are very few publicly traded firms in which an ESOP has more than 20 percent of the firm's stock, and perhaps none in which the plan has a majority of the stock.⁶⁰ Firms in which a majority of the stock is held by an ESOP are, it appears, almost exclusively privately held. And, although the law permits (but does not require) that votes on ESOP stock be passed through to employees in privately held firms, this is done in only a distinct minority of such firms.⁶¹ Thus ESOPs generally do not permit much employee participation in control of the firm in either publicly or privately held firms. In fact, an extensive 1986 survey found that only 4 percent of sampled firms with ESOPs had any nonmanagerial employee representatives on their boards of directors and found no firms with ESOPs in which employee representatives constituted a majority of the board of directors.⁶²

What is particularly interesting is that voting rights have not been passed through to employees even in some firms—including, for most of its young life as an employee-owned firm, the much-publicized Weirton Steel Company⁶³—in which the ESOP owns 100 percent of the firm's stock. Instead, voting rights are commonly held by the ESOP's trustee, who is appointed by a self-perpetuating board. In effect, these firms are operated as nonprofit institutions in which directors with control but no claim on residual earnings are charged with managing the firm as fiduciaries for the benefit of the employees.

A commonly held view is that employee participation in corporate governance is highly desirable but that the risk and the high cost of capital that employees face when they participate in ownership are serious liabilities. From this, one would expect employee ownership to be structured to maximize employees' participation in control but to minimize their contribution of capital. ESOPs, however, have just the opposite character. Since they provide for participation in residual earnings through the purchase of stock (rather than, for example, through a simple profit-sharing plan), they amplify employees' problems of illiquidity and risk bearing. Yet at the same time they typically give employees no voice at all in the management of the enterprise.⁶⁴

If the advantages of employee participation in control exceeded its costs, then one would surely expect to see much more employee control in firms with ESOPs, and particularly in the thousand or more firms in which the ESOP owns a majority of the firm's stock. These firms, or their employees, are already incurring most of the principal costs of employee ownership—particularly the high costs of bearing undiversified risk that employee ownership imposes on employees. The fact that employees typically do not participate in the governance of these firms suggests strongly that the persons responsible for structuring them believe that any reduction in agency costs that might result from making management directly accountable to the firm's employees, even though the employees are already the firm's *beneficial* owners, would be outweighed by the costs—whether in the form of inefficient decisions or high process costs—that would be engendered by the political process required for such accountability.⁶⁵

To be sure, the creation of ESOPs in which votes are not passed through to employees can probably be explained at least in part as an effort by corporate managers to preserve or increase their own autonomy—that is, as a means of protecting the managers both from hostile takeovers and from direct accountability to the employees. But man-

agerial opportunism alone seems insufficient to explain the virtually complete absence of employee control even in firms in which an ESOP holds a majority of the firm's equity. If creation of an ESOP with pass-through of votes to the employees would be more efficient than either no ESOP or an ESOP without pass-through of votes, then one would expect to see hostile takeovers initiated with the aim of creating such a structure—or at least to see such ESOPs put in place upon the successful completion of hostile takeovers attempted for other reasons. In fact, one would expect to see employees themselves, perhaps through their unions, participating in such hostile takeovers.

Beneficial Ownership

As just noted, firms in which an ESOP owns a majority of the company's equity are typically controlled not by their employees but by trustees who manage the firms as fiduciaries for the employees. Thus these firms are not fully employee-owned but rather, in a sense, are nonprofit firms that are operated on behalf of their employees. Indeed, a number of often-cited instances of "employee ownership" in the United States—including, for example, the car rental company Avis, the publisher Norton, and the now-defunct airline People Express⁶⁶—have not, in fact, put voting control of the firm in the hands of their employees; on the contrary, they have given voting rights only to a small group of top managers, leaving most employees to participate only in earnings and not governance.

Similarly, the most successful large-scale experiments with employee-owned industrial enterprise in Britain, the ICOM cooperatives, have also been structured with only beneficial employee ownership. The same is apparently true to a substantial degree of many French and Italian worker cooperatives, and one can even argue that this characterization has some application to Mondragon. In short, examples of successful employee ownership in industrial enterprise frequently involve not true worker control but rather firms that are only beneficially owned by their workers, with actual authority exercised by fiduciaries.

A plausible inference to be drawn from this pattern, as already noted, is that full employee control tends to bring serious inefficiencies with a work force as heterogeneous as those to be found in most industrial firms. One might concede this point, however, and nevertheless argue

that even firms that are only beneficially owned by their employees are often preferable to investor-owned firms. Indeed, many advocates of employee self-governance explicitly call for giving employees only attenuated property rights in their firms' earnings and assets.⁶⁷

Some of the most important benefits of employee ownership discussed in Chapter 5 could presumably still be realized with only beneficial ownership instead of full employee ownership. The problems of monitoring employees, employee lock-in, strategic bargaining behavior, and poor communication of employee preferences arise principally because of the conflict of interest between employees and owners that characterizes an investor-owned firm. Even beneficial employee ownership promises to eliminate that conflict of interest in substantial part.

At the same time, serious countervailing inefficiencies may result if industrial firms are structured, to an important degree, as nonprofits in which employees have beneficial ownership but lack essential elements of full ownership, such as ultimate authority to replace top management or to compel distributions of earnings or assets. We will explore the nonprofit form more carefully in Chapter 12, and here simply note that such a structure largely insulates the firm's management from any direct incentives to perform efficiently. Further, such firms are subject to special problems in adjusting their capital stock. When nonprofit firms need a rapid infusion of capital, they often have difficulty raising it since they cannot sell equity shares. Conversely, when market conditions permit, nonprofit firms tend, by retaining earnings, to accumulate capital well beyond the amount needed for efficient operation.

For these reasons, only under two conditions does it make sense to organize an industrial firm as a nonprofit, managed in substantial part on behalf of but not by its employees: (1) if severe inefficiencies in contracting for labor would result if the firm were investor-owned; and (2) if collective governance by the employees would be extremely costly. In such circumstances, the inefficiencies of the nonprofit form might be smaller than either the costs of labor market contracting that would accompany investor ownership or the costs of collective decision making that would accompany true employee ownership.

The fact that employee ownership in the industrial sector is confined largely to firms that are only beneficially owned by their employees rather than being fully employee-controlled suggests that condition (2) holds. The fact that even beneficial employee ownership rarely appears in complex industrial firms where it is not heavily subsidized suggests

that condition (1) does not hold. Taken together, these conclusions imply two things. First, if employees are to be given some form of ownership interest in an industrial firm, it is most efficient to structure the firm as a fiduciary entity managed on behalf of, but not fully controlled by, its employees. Second, the costs associated with contracting for labor on the market—that is, the inefficiencies of investor ownership—are nevertheless insufficient to justify the costs of adopting the nonprofit form to avoid them. As a general matter, industrial firms that are only beneficially owned by their employees may be more efficient than fully employee-owned firms, but are evidently less efficient than investor-owned firms.

It should be kept in mind that we are talking here only of what we have loosely termed “industrial” firms—that is, firms with sufficient hierarchy and division of labor to have a work force with significantly heterogeneous interests. As we saw in Chapter 5, employee ownership can be, and is, quite successful when the employee-owners can be drawn from a relatively homogeneous group for whom the costs of collective decision making are not an important issue.

Codetermination

German-style codetermination, providing for substantial worker participation in corporate governance but no direct participation in net earnings, has roughly the opposite characteristics from ESOPs.⁶⁸ More particularly, under legislation adopted for the coal and steel industry in 1951 and extended to German industry in general by stages in 1952 and 1976, German workers are entitled by law to elect half of the members of the supervisory board of directors in all large German firms, though workers otherwise remain salaried employees with no equity stake in the firm.⁶⁹

If all workers were homogeneous, and if the workers had complete parity on the corporation’s board, then codetermination would yield a structure roughly like a two-person partnership, in which the two parties must continually bargain between themselves over all relevant matters concerning the operation of the firm. It would also have something of the character of ordinary collective bargaining, except that it would have the advantage that information would be fully shared and that the terms of the agreement might be more easily adjusted as time goes along.⁷⁰

Such a structure is, arguably, roughly what one finds in the German

coal and steel industry. The worker representatives there have been granted substantial parity on the board of directors. Moreover, those worker representatives arguably speak with a roughly unitary voice as a consequence of the heavily blue-collar nature of the work force and, probably much more important, of the substantial role given to the union in representing the workers on the board. The highly centralized structure of union authority in Germany presumably tends to yield a common policy for the worker representatives vis-à-vis the investor representatives on the boards in these firms.

The industries covered by the Co-Determination Act of 1976 (namely, all firms outside of coal and steel in which there are more than 2,000 employees) present a somewhat different picture. These firms evidently exhibit all the work-force heterogeneity typically encountered in modern industry. And the unions have not been given a central role in the process of selecting worker representatives to the supervisory board. As a result, the worker representatives on the board represent constituencies with diverse interests. The legally mandated system for selecting worker representatives reinforces this, because it requires that there be included at least one representative from each of three classes of workers: wage earners, salaried employees, and managerial employees. From all that has been said above, one would not expect that this system of representation would be highly viable as a means of governing the firm. Given the apparent difficulty of making collective self-governance workable for employees alone when the labor force is heterogeneous, it would be surprising if a firm’s electoral mechanisms, including voting for and within the board of directors, could effectively be employed not only to resolve conflicts among different groups of employees but also to deal with the more serious conflicts of interest between labor and capital.

The German experience does not clearly belie that expectation. To begin with, codetermination has been imposed upon German firms by force of law; no similar system seems to have been adopted by any significant number of firms either inside or outside of Germany in the absence of compulsion.

In addition, a variety of elements in the German system of codetermination seem well designed to prevent politics at the board level, both among workers and between workers and shareholders, from yielding instability and inefficient decisions. First, outside the coal and steel industry, labor lacks full parity on the board: if there is deadlock between the worker and the shareholder representatives, a tie-breaking

vote is effectively granted the shareholder representatives. Second, German firms have a dual board structure in which the higher "supervisory" board, on which the worker representatives sit, is distant from all but the broadest decisions of policy. Third, the provision for separate managerial representatives among the employee representatives will often side with the shareholder representatives. Fourth, the matters of most direct importance to workers are not decided at the level of the supervisory board. Either they are decided higher up, at the industry level, through collective bargaining between the unions and the employers' associations, or they are decided at a lower level in dealings between management and the legally mandated works councils. As a consequence, the board is not used for making the particular types of decisions where workers and shareholders have the most strongly conflicting interests, and where the workers have the most strongly conflicting interests among themselves.

Indeed, it appears that codetermination has not had a substantial impact on firm decision making at the board level, which continues to be dominated by shareholder interests in firms outside the iron and steel industry. The worker representatives to the supervisory board arguably play a largely informational role, providing a credible source of information from the firm to the workers (and vice versa) in support of collective bargaining by the unions and decision making in the works councils, where the real worker influence takes place and where shared information presumably reduces the incentive for, and hence the costs of, strategic bargaining. This suggests, in turn, that there may not be much difference between the regime imposed on those firms covered by the 1976 law and the smaller firms still covered by the 1952 law, which provides for only one-third of the directors to be worker representatives, and indeed that both these systems may be little different in practice from the regimes imposed in the Scandinavian countries and elsewhere, under which workers have the right to have only one to three members on the board—enough to serve as informational conduits, but not enough to exercise meaningful control.

Unionization

In those firms in which employees bargain collectively through labor unions, employees already employ a political process for the purpose of aggregating their individual interests. The difference between such an

arrangement and the type of collective representation involved in employee ownership, of course, is that the union's political process is used not to select the firm's management but to select representatives to bargain with a management chosen by the firm's shareholders.

It might seem that unions have most of the costs and few of the benefits of employee ownership. On the one hand, because unions do not involve full employee ownership, they do not entirely remove the possibility that the management of the firm will behave opportunistically toward the employees (or vice versa). Yet, on the other hand, they potentially have all the costs of collective decision making among employees.

There is surely some truth to this view, and it presumably helps explain the near abandonment in the contemporary United States of the adversarial model of collective bargaining that was adopted in American law in the 1930s.⁷¹ Whatever the overall efficiency of that model of employee representation, however, there are many ways in which it adapted to the problems of collective representation on which we have been focusing.

For one thing, employees with managerial or supervisory responsibilities are generally not unionized; it is usually only the employees who make up the lower, more horizontal strata among the firm's employees who belong to a union. (Where more senior employees form unions, they tend to be highly homogeneous, as in the case of teachers or pilots.) Where the jobs held by the unionized employees are particularly diverse, the employees are frequently split up into separate bargaining units. As a consequence, there is commonly a fair degree of homogeneity of interest among the employees represented by any given union.⁷²

Further, unions typically bargain with management only over a relatively narrow range of issues immediately touching on the employees' interests, such as wages, hours, and job classifications.⁷³ Other issues, such as the firm's investment policies or even its policy on layoffs, are seldom bargained over even though in theory it might be more efficient if employees were more actively involved in deciding them.⁷⁴ Indeed, unions themselves seem to avoid broader involvement of this sort, intentionally keeping the scope of bargaining narrowly confined.⁷⁵ There may be a variety of reasons for this. But whether it is cause or consequence, with this strategy a union averts certain possibilities for costly internal conflict. By confining themselves to such matters as wages, hours, and job classifications, unions avoid the ne-

cessity of making a variety of hard choices where the interests of their members conflict. They leave that task to management and stay focused on the less controversial aim of generally pressing for greater equality with respect to the subjects they bargain over.

Finally, it is often observed that unions are seldom democratic.⁷⁶ This is commonly deployed in both the social science and the policy literature, much as the general absence of genuine shareholder democracy in publicly held business corporations was deplored several decades ago. But it may be that greater democracy would bring much higher governance costs without a corresponding improvement in the accuracy with which the union members' preferences are represented. Michels's iron law of oligarchy⁷⁷ may in fact be an economic law, at least where unions are concerned.

Similar considerations may help explain why it is that bargaining between a union and a firm is often conducted in large part by representatives from the union's national office and not just by local union officials: this helps defuse even further the problem of local internal politics.

Experience with collective bargaining thus does not provide strong evidence that democratic processes can be effectively employed to represent the interests of a heterogeneous class of employees in general corporate decision making.

Other Problems with Employee Governance

One might concede that the high costs of employee governance have been a critical constraint on employee ownership but nevertheless believe that these costs have some source other than conflicts of interest among the employees or that, whatever the source of the costs, they can be made manageable through experience and organizational innovation. Two such arguments deserve special attention.

Employees May Lack Managerial Skills

It is sometimes argued that employee ownership is common among service professionals but rare among industrial employees simply because the latter lack the skills necessary for governance of a firm. For example, blue-collar employees may have insufficient knowledge of management or finance to select or police the firm's managers effectively. Or such

employees may be inclined to be short-sighted in planning; the high salience of wages and working conditions may make the employees focus on those concerns to the detriment of new investment.

Perhaps there is some truth in this view. The prevailing patterns of ownership suggest, however, that the skill level of employees is much less important than their homogeneity of interest in making employee ownership feasible. Employee ownership is evidently viable among blue-collar employees where there is little diversity of interest, as with the drivers in the transportation cooperatives or the semi-skilled employees in the plywood manufacturing cooperatives. Conversely, employee ownership is rare among white-collar employees who are not highly homogeneous, such as those employed in retailing, hotel and restaurant services, health care other than physician services, or computer programming.

An individual employee need not herself have the expertise to make managerial decisions in order to exercise her voice effectively as an owner. She need only be able to vote intelligently in electing the firm's directors. Consider, for example, a firm like General Motors. It is quite believable that even an assembly line worker at GM is in a position to act more thoughtfully in electing directors than are most of that firm's public shareholders, since she is likely to be more willing and able to obtain relevant information about the candidates and to act on it. Furthermore, many of GM's employees are not assembly line employees but financial planners, design engineers, and marketing executives—employees who have a great deal of information relevant to an assessment of the firm's management and who are likely to influence less-informed blue-collar employees in deciding whom to elect as directors. Consequently, it is quite plausible that, if the employees owned General Motors, they would elect a more effective board of directors and hold them more accountable than do the firm's current public shareholders—or at least it is quite plausible that this would be the case if, contrary to fact, there were no important conflicts of interest among the firm's employees.

Employees Lack Experience with Governance

It has also been argued that a major obstacle to widespread employee ownership is simply the absence of the customs, mores, and standard procedures necessary to make employee governance effective.⁷⁸ Ac-

cording to this theory, employees must first become accustomed to the notion of managing the firm where they work and develop the experience and methods needed for the task. This is presumably a cumulative process; employees in one firm can benefit from the example of another. Once the proper institutions and procedures are well established and there is substantial accumulated experience in working with them, employee ownership will compete effectively with investor ownership in a broad range of settings, including those in which there is substantial heterogeneity of interest among the employees.

We cannot easily dismiss this argument. Organizational innovation and its diffusion have sometimes proceeded slowly in other settings. But there is strong reason to be skeptical.

Institutions for collective governance play a central role in American culture and are familiar to most citizens in a broad range of settings, from presidential elections to the student council and the local Moose Lodge. Employees, in particular, have long been familiar with collective governance in unions. In such a cultural environment, it is hard to argue that lack of experience with collective decision making in general is a major obstacle to the viability of employee ownership. At most, one can argue that the obstacle is lack of knowledge about, or experience with, the specific types of governance mechanisms that are needed to permit employees to act collectively in managing industrial enterprise.

The experience with ESOPs and at Mondragon suggests that the most suitable forms for employee self-governance in large-scale enterprise may involve a complex combination of representative and fiduciary mechanisms. Further experimentation and experience in developing these organizational forms may therefore help to reduce their costs. Perhaps there are sufficient gains yet to be had in this respect to yield important improvements in the viability of employee ownership. But this seems unlikely. Even in the United States there has already been considerable experience with forms of employee self-governance in service enterprise and in small-scale industry. Further, as later chapters will discuss, there has accumulated over many decades substantial experience with consumer cooperatives and with other types of producer cooperatives (such as farmer-owned processing and marketing cooperatives), some of which rank among the nation's largest industrial firms. These firms provide obvious models for the governance of worker cooperatives. Their success suggests strongly that where there are substantial gains to be had from collective ownership,

its development will not long be inhibited simply by lack of experience with institutions for collective governance in the particular setting involved.

The experience with housing cooperatives and condominiums, which is explored in detail in Chapter 11, is instructive in this regard. Prior to 1961 the condominium form was essentially unknown in the United States, and its close cousin, the housing cooperative, was largely confined to the wealthy. Since then, however, changes in property law have made the condominium form feasible, and developments in taxation have made it financially attractive for a broad range of individuals. As a consequence, the condominium form has now spread rapidly and widely through the housing market even though it requires, at its core, a mechanism for collective self-governance by the owner-occupants. There is reason to believe that collective governance in housing condominiums is costly and that, in part for that reason, without substantial tax subsidies the condominium form would be far less common than it is. It also appears that the mechanisms for collective governance in condominiums are continuing to evolve—for example, by delegating increasing authority to hired management—in ways that reduce these costs. Nevertheless, the basic governance mechanisms necessary to make condominiums viable in competition with investor-owned (that is, rental) buildings developed quite rapidly. If so many individuals can so quickly become accustomed to collectively governing the apartment building where they live, why can they not also quickly become accustomed to collectively governing the firm where they work? Evidently more is required to make employee self-governance viable than simply additional time to get accustomed to it.

A Test Case: United Air Lines

The employee buyout of United Air Lines in 1994 promises to provide an illuminating test of the importance of governance problems and of the potential mechanisms to cope with those problems. Originally the company's 7,000 pilots sought to purchase the firm by themselves, and in fact reached a preliminary agreement with the company's shareholders for such a transaction in 1989. If that transaction had succeeded, the result would have been very much the type of employee-owned firm that has proved successful elsewhere, with ownership shared by a single highly homogeneous group of employees. The pi-

lots' 1989 buyout ultimately failed, however, in part because of the strong opposition of United's 23,000-member machinists' union, reflecting the serious conflicts of interest between these two quite different groups of employees.

In the successful 1994 buyout, the pilots brought the machinists in as participants. The participating employees acquired, through ES-OPs, 53 percent of United's stock, with the remaining 47 percent left in the hands of public shareholders. The consequence is a conspicuously heterogeneous set of owners, with potential conflicts of interest not only between the employees and the public shareholders but also, among the employees, between the pilots and the machinists. In addition, the machinists are themselves, in comparison with the pilots, a highly heterogeneous group. In an apparent effort to cope with these potential conflicts, United's governance mechanisms were revised, as part of the buyout, to provide for selection of the company's board of directors in a complex and highly constrained fashion that involves no direct voting for any of the board members by either the pilots or the machinists. Rather, majority control of the board is in the hands of quasi-self-appointing outside directors who are apparently intended to serve as fiduciaries simultaneously for all of the factions among the owners.⁷⁹ (In addition, both the pilots' and machinists' unions committed themselves, as one of the terms of the buyout, to six-year no-strike agreements.)

United's new governance structure reflects the general pattern we have seen elsewhere, with substantial muting of electoral democracy and heavy reliance on control by fiduciaries. The size, economic importance, and potential profitability of the airline make this an important experiment. It will be interesting to see whether this complex ownership structure proves viable in the long run, suggesting that participation in ownership is a promising route for restructuring labor relations across the economy in general, or whether it ends up being simply a temporary expedient to resolve the bargaining impasse reached under the old regime of adversarial collective bargaining, with the company ultimately reverting to a more homogeneous class of patron-owners.

Conclusion

The classical model of the business firm, under which formal control is confined to suppliers of capital while management in turn deals with

employees through market contracting, leaves room for considerable inefficiency in the form of agency costs between owners and managers and opportunistic or strategic behavior between the firm and its employees. In theory, employee ownership promises substantial efficiency improvements in these respects. And in practice it appears that, when the employees involved are highly homogeneous, employee ownership often is more efficient than investor ownership. The evidence suggests, however, that with a heterogeneous work force direct employee control of the firm brings substantial costs—costs that are generally large enough to outweigh the benefits that employee ownership otherwise offers. For similar reasons, true sharing of control between labor and capital does not appear promising as a route to efficiency. Paradoxically, the aspect of employee ownership often extolled as its principal virtue—active participation in governance of the firm through democratic institutions—appears in fact to be its greatest liability.

Consequently, forms of employee participation that fall short of true ownership may offer better prospects for improving on the efficiency of the classical model. For example, simple profit-sharing plans may be a better approach to increasing incentives for productivity, while quality circles, shop floor committees, works councils, labor-management committees, and informational seats for labor on the board of directors may be more workable means of improving the flow of information between management and employees.

To be sure, all such conclusions must be tentative. Organizational innovations may yet make employee ownership viable in circumstances where it has previously failed to make headway. But for the present, it seems most reasonable to predict that successful instances of employee ownership will remain largely confined to firms with highly homogeneous classes of employee-owners.

The chapters that follow, which explore other forms of producer and consumer-owned enterprise, strongly reinforce the conclusion that the problems of collective decision making play a key role in determining patterns of ownership. Indeed, that role is often much easier (and less controversial) to identify in other forms of shared ownership than it is in employee ownership.