In some situations, however, it is infeasible for the firm to own its patrons. In particular, when the patrons are individuals such as workers or consumers, legal prohibitions on personal servitude, as well as a variety of practical contracting problems, obviously bar this arrangement. If the firm and its patrons are to be connected by ownership, the patrons must own the firm.

For related reasons, ownership of the patrons by the firm can sometimes be impractical even where the patrons are not individuals but instead are other firms. Consider the common case—discussed at length in Chapter 8—of a wholesaler owned as a cooperative by the retail stores to which it sells. The problems of market failure to which this ownership arrangement responds (typically market power on the part of the wholesaler) might alternatively be solved by having the wholesaler own the retail stores. And, of course, fully integrated chain store operations of the latter type are common. But that arrangement can create diseconomies of scale, including loss of the strong incentives for efficient operation that exist when the individual retail stores are owned separately by their local managers. Having the stores collectively own their supplier, rather than vice versa, can be the superior arrangement. In short, the costs of ownership are often asymmetric between a firm and its patrons—a point that emerges even more clearly in the next chapter.

3

The Costs of Ownership

We have observed that ownership has two essential attributes: exercise of control and receipt of residual earnings. There are costs inherent in each of these attributes. Those costs fall conveniently into three broad categories: the costs of controlling managers, the costs of collective decision making, and the costs of risk bearing. The first two categories are associated with the exercise of control. The third is associated with the receipt of residual earnings. All of these costs can vary substantially in magnitude from one class of patrons to another.

We shall survey these three types of costs here in general terms. As with the costs of market contracting surveyed in the preceding chapter, subsequent chapters will offer deeper analysis and more copious and detailed illustrations.

Costs of Controlling Managers

In large firms, and especially in firms with a populous class of owners, the owners must generally delegate substantial authority to hired managers.[1] Thus, in widely held business corporations, as in large cooperatives, most decision-making authority is delegated to the firm's board of directors, who in turn delegate most operational decisions to the firm's senior officers. This delegation brings with it the costs commonly labeled "agency costs." For our purposes, these costs can conveniently be broken down into two types: the costs of monitoring the managers and the costs of the managerial opportunism that results from the failure to monitor managers with perfect effectiveness.[2]

Monitoring

If the patron-owners of a firm are to control its management effectively, they must incur the costs of (1) informing themselves about the operations of the firm, (2) communicating among themselves for the purpose of exchanging information and making decisions, and (3) bringing their decisions to bear on the firm's management. I shall refer to these costs collectively as "monitoring costs." These costs can vary substantially among different classes of patrons. Since patrons are likely to accumulate information about the firm simply as a by-product of transacting with it, the cost of monitoring for a given class of patrons will generally be inversely proportional to the importance, frequency, and duration of the patron's transactions with the firm.[3] The costs of monitoring will also depend on the ease of organizing the patrons for collective action, which may depend in turn on factors such as the patrons' physical proximity to one another and to the firm.

For example, tenants in an apartment building generally have relatively low monitoring costs. They deal repeatedly with the building's management, often for a number of years, in transactions that involve a significant fraction of their budget. They therefore have both the opportunity and the incentive to learn a great deal about how well the building is managed. Close proximity also permits easy organization for collective action. These are important factors in the viability of tenant ownership of apartment buildings through cooperatives and condominiums, as will be discussed further in Chapter 12.

Finally, the number of patrons among whom ownership is shared affects monitoring costs. If all patrons are to participate effectively in decision making, then a large class of owners requires substantial duplication of effort in becoming informed. Moreover, the monitoring efforts of any individual owner have the properties of a public good for the owners as a group: the benefits of that monitoring are enjoyed by all other owners as well, regardless of whether they have undertaken any monitoring of their own. Consequently, as the number of owners grows, each individual owner's share of the potential gains from effective monitoring decreases, thus reducing the individual's incentive to monitor.

It follows that, where the class of owners is large, it may be prohibitively costly to induce the owners to undertake anything beyond the most cursory monitoring. In itself, this argues for the smallest group of

owners possible—preferably a single owner. The fact that, despite this, a large firm often has a very large class of owners therefore suggests that either or both of two things must be true. First, the costs of market contracting would be much higher under any alternative assignment of ownership. Second, the costs of managerial opportunism are modest even though the firm's owners cannot actively supervise the managers. We shall first explore the latter possibility. Then, at the end of the chapter, we shall return to the former.

Managerial Opportunism

To the extent that the owners of a firm fail to exercise effective control over its managers, the managers have an opportunity to malinger or engage in self-dealing transactions. Clearly this can sometimes be costly.[4] Yet the conduct of a firm's managers is conditioned by a variety of constraints and incentives beyond direct sanctions or rewards from the firm's owners. There are important limits to the costs of managerial opportunism even in firms whose nominal owners are in a poor position to do any active monitoring of the firm's management at all.

Consider first self-dealing. The transactions necessary for managers to divert to themselves a significant fraction of the residual earnings in a large firm are often difficult to conceal. Moreover, these transactions are in most cases explicitly proscribed by contract or by law, thus exposing the managers to a variety of moral, contractual, tort, and criminal sanctions that can be brought to bear without collective action on the part of the firm's owners. In particular, self-dealing managers expose themselves to shaming by fellow workers, friends, or family, to derivative suits initiated by individual shareholders or enterprising lawyers, and to civil or criminal prosecution by the state (including, conspicuously, the tax authorities).

To be sure, although legal, contractual, and moral constraints may generally suffice to keep managers from putting their hand in the till, they will not necessarily ensure that managers work hard and make effective decisions. Again, however, pride and moral suasion provide important motivation, particularly for the types of individuals who work their way to the top of a managerial hierarchy. The need for the firm to prosper if managers are to keep their jobs or, even better, to enhance them, also provides an important work incentive.[5] Moreover, it may be a mistake to exaggerate the degree of effort or ingenuity that is required of the senior managers in a typical business enterprise, and thus the potential gains from better monitoring of those managers by a firm's owners. In many firms, imitation of standard managerial practices may suffice for relatively successful performance.

In sum, the inability of a firm's nominal owners to exercise much direct control may result in only a modest amount of organizational slack, at least when compared with any realistic alternative.[6] Indeed, in the chapters that follow we shall encounter large groups of firms (including mutual life insurance companies and nonprofit hospitals) that have been successful over long periods of time in competitive environments without any effective exercise of control by owners whatever—often without even having any owners.

There is, however, one costly managerial perquisite—excessive retention of earnings—that is not easy to detect or proscribe, that is likely to bring approval rather than censure from friends and colleagues both inside and outside the firm, and that is generally encouraged rather than checked by managers' desires to retain or build their empire. Retentions benefit managers by creating a buffer against adversity and by increasing the size of the firm that the managers control. But retentions are costly to the firm's owners if the rate of return on the retentions is less than the return available on investments outside the firm or if, regardless of the rate of return the retentions bring, the funds retained can never be recovered by the current owners (as happens in some mutuals and cooperatives). This problem is most easily discerned in nonprofit[7] and mutual firms, but it is arguably the principal source of inefficiency in investor-owned firms as well.[8] And because excessive retention of earnings tends to enhance rather than decrease the survival value of a firm, those firms that are particularly subject to this tendency—as firms with diffuse ownership are—may actually be favored rather than pressured by the invisible hand of market selection.

Whatever the nature of the managerial opportunism involved, where the losses it brings are smaller than the costs of the monitoring that would be required to prevent it, it is of course efficient for the firm's owners to tolerate the opportunism. Agency costs, therefore, are the sum of the costs incurred in monitoring and the costs of managerial opportunism that result from the failure or inability to monitor with complete effectiveness.

Collective Decision Making

When many persons share ownership of a firm, there are likely to be differences of opinion concerning the firm's policies and programs. Sometimes those differences will merely reflect different judgments about the most effective means for achieving a shared goal. More serious differences arise, however, when the outcome of the decision will affect different owners differently. Broadly speaking, this could happen for either of two reasons.

First, the individuals involved may differ in the way in which they transact with the firm as patrons—that is, in the nature of the goods or services they sell to, or purchase from, the firm. To take a simple example, a decision to repair the elevators in a four-story cooperative apartment building will benefit the first-floor residents much less than those on the fourth floor. The residents, depending on where they live in the building, may therefore disagree on the desirability of paying costly overtime to get the repairs done quickly. Similarly, if a worker-owned firm must shut down one of its two plants, the workers at the two plants are likely to have very different preferences about which plant should be chosen.

Second, the owners may have differences in preferences that arise from their personal circumstances rather than from any differences in their transactions with the firm. A decision by a cooperative apartment building to accelerate repayment of the principal on the building's mortgage may affect members differently depending on their personal liquidity and tax status even if they occupy identical apartments and have identical leases. Or a decision by a worker-owned firm to shift to riskier lines of business, and thereby increase the chance that the firm will fail, is likely to be less attractive to older workers than it is to younger workers who, though doing the same job, are more easily retrainable and have fewer ties to the local community.

In order for a firm's owners to make decisions when their interests differ, they must employ some form of collective choice mechanism. The nearly universal approach is to adopt a voting scheme, with votes apportioned either by volume of patronage or on the basis of one-member-one-vote. When the interests of the individual owners are diverse, such mechanisms for collective choice engender costs. These costs, which for future reference we can label the "costs of collective decision making," are logically distinct from agency costs. They can be large even in firms, such as modest-sized partnerships, in which there are no hired managers and hence no significant agency costs. Conversely, the costs of collective decision making can be negligible in large corporations in which ownership is widely shared and hence agency costs are large, as long as the owners have highly homogeneous interests.

To make this distinction clear, we can define "agency costs" as the costs of monitoring and managerial opportunism that the firm would incur even if the interests of all owners were identical. The "costs of collective decision making" are then the additional costs that result from heterogeneity of interests among the owners. Unlike agency costs, the costs of collective decision making have been largely neglected in the literature on corporate control and the economics of organizational form.[9] Nevertheless these costs play a crucial role in determining the efficiency of alternative assignments of ownership.

The collective choice mechanisms employed within firms are essentially political mechanisms. Their costs are therefore characteristically the costs of political mechanisms in general. In recent decades, the "public choice" literature has begun to provide a more systematic understanding of these costs, which might be termed the costs of "political failure," analogous to the costs of "market failure" that affect market mechanisms. Although that literature still leaves us with a very partial understanding of these costs, some general characterizations are possible.

The costs associated with collective choice mechanisms are of two broad types. First, there are the costs resulting from inefficient decisions—that is, from decisions whose outcomes fail to maximize the aggregate welfare, or surplus, of the owners themselves as a group. Second, there are the costs of the decision-making process itself.

Costly Decisions

Inefficient decisions can arise in several ways. To begin with, as already noted, majority voting tends to select the outcome preferred by the median member of the group, while efficiency generally calls for the outcome preferred by the average member. Where the median and the average member have substantially different preferences, voting can produce seriously inefficient decisions.[10] Consider again the hypothetical four-story cooperative apartment building with a broken elevator. If the residents of the first two floors, who do not use the elevator, outnumber the residents of the top two floors who do, then the residents as a whole might vote not to pay overtime to hasten the repairs, even though the money thus saved is substantially less than the costs, both pecuniary and nonpecuniary, that the delay imposes on the residents of the upper floors.

Alternatively, control over the political process can fall into the hands of an unrepresentative minority who, intentionally or unintentionally, use that control to make decisions that inefficiently exploit the majority in favor of the minority. This is particularly likely to happen when, as is often the case, some patrons are better situated to participate effectively in collective decision making than others—perhaps because they have few other demands on their time, or have special managerial expertise, or have special access to information. For example, governance of a cooperative apartment building might be dominated by those residents of the building who are retired, even if they are in the minority, because they have more time to attend meetings. As a consequence, improvements that primarily benefit the retirees, such as elevator repairs, might be emphasized at the expense of those that do not, such as repairs to the children's playground, even if the reverse priorities would be more beneficial to the building's occupants as a whole.

Whether it is the majority that inefficiently exploits the minority or vice versa, the dominant group need not be particularly venal for the resulting costs to be substantial. It is sufficient that, as is natural, the decision makers' own interests simply have more salience for them than do the interests of others.

Costly Process

The costs of the collective choice process, in turn, may also have several sources. Even if individual owners always seek to exercise their right of control without opportunism and to reach the decisions that will be most efficient for the owners as a whole, they may need to invest considerable time and effort in obtaining knowledge about the firm and about other owners' preferences, and in attending the meetings and other activities necessary to reach and implement effective collective decisions. We also know from public choice theory that the possibility of a voting cycle[11] among alternatives increases as preferences among the electorate become more heterogeneous.[12] Such cycling may be costly if there are transaction costs involved in repeatedly altering the firm's policies. More important, the instability that underlies cycling can give extraordinary power to those in control of the voting agenda to obtain the outcomes they desire, no matter how inefficient those outcomes may be.[13] Finally, if owners seek to behave strategically, then further costs may result from efforts to hide or discover information or to make or break coalitions.

Methods exist for limiting these process costs. Delegation of authority to committees, for example, can reduce the costs of participation, inhibit cycling, and facilitate vote trading that will mitigate the median voter problem. But delegation can also produce seriously inefficient outcomes by empowering committee members to impose their own idiosyncratic preferences on the group as a whole.[14]

Resolving Conflicts

Even if the owners of a firm are heterogeneous in their interests, the costs of collective decision making may nevertheless be low if there is some simple and salient criterion for balancing those interests. Consider the division of the firm's net earnings among its owners. This is potentially controversial where the character or volume of the transactions between individual owners and the firm varies substantially. Important examples, which we shall examine closely in Chapter 6, involve employee-owned firms in which the employees differ in the types of work they do. The costs of reaching agreement on an allocation of earnings, and the possibility that the resulting allocation will create inefficient incentives, may be manageable if it is easy to account separately for the net benefits bestowed on the firm by transactions with individual owners and to apportion the firm's earnings according to that accounting. Alternatively, if the value of each individual owner's transactions with the firm is difficult to measure, a rule of equal division may serve as a focal point[15] on which agreement can easily be reached, thus minimizing the process costs of decision making though perhaps creating some inefficient incentives. Law firms often follow one or the other of these approaches: some use explicit multifactor productivity formulas to determine partners' shares; others follow a simple rule of equal division of earnings among all partners of a given age. Where such clear and conventional decision-making criteria are

absent, however, workable agreement among the owners can take a long time to reach, and may in fact never be reached.[16]

Participation

In some cases, the process of collective decision making arguably yields benefits for the patrons involved and not just costs. In fact, advocates of worker ownership often suggest that participation in control of the firm through democratic processes is of value in itself, quite apart from the practical import of the substantive decisions that result,[17] and a similar argument is sometimes made on behalf of consumer cooperatives and other forms of noncapitalist enterprise.[18] Although the reasons for valuing participation in this way are seldom spelled out explicitly, at least three can be identified.

First, individuals might simply enjoy the experience of participating in collective decision making—attending meetings, debating alternatives, assuming offices—as a social activity that is satisfying in itself. That is, political activity may in effect be a consumption good. Second, as is sometimes argued in the context of worker ownership, individuals may gain psychological satisfaction from the feeling of being in control, and this feeling may be enhanced for a firm's patrons by permitting them to participate directly in the decision making of the firm.[19] Third, as has also been argued on behalf of worker ownership in particular, participation in collective decision making within the firm may be useful training for participation in the democratic political processes of the larger society, and might be valued for this reason not only by the individuals involved but also by the rest of society.[20]

But note that these benefits, real though they may be, still involve tradeoffs. To grant the franchise and the associated benefits of participation to one group of patrons typically requires denying them to all other groups of patrons. Advocates of alternative forms of ownership sometimes overlook this point. For example, it has been argued, on behalf of worker ownership, that it is inconsistent to have democracy at the level of the state and not at the level of the firm.[21] Yet in fact there is democracy in the typical investor-owned firm; it is just that the investors of capital do the voting rather than the workers. Converting to worker ownership means not only enfranchising the workers but also disenfranchising the firm's investors while continuing to deny the franchise to the firm's consumers. Consequently, the question generally is not whether there is voting in a firm, but rather who votes. If the benefits of participation as a good in itself are greater for one group of the firm's patrons than for another, then this becomes a further consideration in assigning ownership.

The value to individuals of participation as a good in itself is an empirical question that is illuminated by the analysis of existing ownership patterns in subsequent chapters. Interestingly, the evidence suggests strongly that for all classes of patrons—including, in particular, employees—the benefits of participation are generally insufficient to outweigh the costs of collective decision making.

Why Not Make Everybody an Owner?

In theory it would be possible to have all classes of patrons share in collective decision making, and thus not completely disenfranchise anyone. This is essentially the position taken by those who feel that every group affected by a business firm's decisions—its "stakeholders," such as workers, customers, suppliers, members of the local community, and environmental groups—should have representation on the firm's board of directors.[22] Moreover, one might think that this would also have the important advantage of reducing the costs of market contracting for all of the firm's patrons and not just for a single group of them.

But because the participants are likely to have radically diverging interests, making everybody an owner threatens to increase the costs of collective decision making enormously. Indeed, one of the strongest indications of the high costs of collective decision making is the nearly complete absence of large firms in which ownership is shared among two or more different types of patrons, such as customers and suppliers or investors and workers.

Risk Bearing

The preceding discussion has focused on the costs associated with the first element of ownership: the exercise of control. But there are also costs associated with the second element of ownership: the right to residual earnings. Most conspicuous among these is the cost of bearing important risks associated with the enterprise, since those risks are often reflected in the firm's residual earnings.[23] One class of a firm's patrons may be in a much better position than others to bear those

risks—for example, through diversification. Assigning ownership to that class of patrons can then bring important economies.

This is a familiar explanation for the prevalence of investor-owned firms. It is not true, however, that lenders of capital are the only low-cost risk bearers. For example, customers can also be in a good position to bear the risks of enterprise, particularly where the goods or services involved are a small fraction of the customers' budget or where the customers are themselves firms that can pass the risk on to their own owners or customers. Moreover, the existing literature often imputes to a firm's noninvestor patrons, and to employees in particular, a greater degree of risk aversion than they actually seem to exhibit. Indeed, the evidence offered here suggests that the importance of risk bearing as an explanation of ownership is commonly overstated.

Entrepreneurship

So far we have been focusing on the costs of ownership for an established firm. But there are also costs associated with organizing a firm in the first place or with changing a firm's form of ownership. We can think of these costs as the costs of entrepreneurship.

If, initially, the prospective owners of a new firm had to assemble and organize themselves on their own before establishing the firm, then it would generally be impossible for any numerous and widely dispersed class of patrons to assume ownership. But in fact the organization of a firm is generally brokered. An entrepreneur first establishes the firm by herself and then sells it to the patrons who will ultimately own it. In the process, the entrepreneur organizes the patrons into a group.

For example, widely held business corporations are typically organized first as closely held firms. Subsequently, shares are sold off to members of the investing public in a stock offering brokered by an investment banking firm. Similarly, new condominium and cooperative housing is usually built by a single developer who initially owns the entire building and then sells the separate units to individuals who ultimately become, collectively, the owners of the building. And the numerous worker-owned plywood manufacturing cooperatives in the Pacific Northwest, discussed in Chapters 5 and 6, were in many cases established by individual promoters who would form a company and then find workers to buy it.

Established firms, moreover, can often change their form of ownership relatively easily. For instance, over the past century a number of investor-owned insurance companies have converted into mutual (policyholder-owned) companies and vice versa. Since the 1970s, large numbers of apartment buildings have converted from investor ownership (that is, rental) into cooperatives or condominiums. And more recently a number of investor-owned industrial firms have been sold to their workers. Because such transactions can be brokered, the costs of the transactions are often modest relative to the value of the firm. As a consequence, the costs of changing forms of ownership need not have an important bearing on the forms that ultimately survive. Two factors can, however, make the costs of changing a serious impediment.

First, important economies derive both from the presence of established brokers who specialize in ownership transactions and from the existence of standardized procedures for handling those transactions. Where such institutions have not yet developed, the costs of adopting or converting to a particular form of ownership may be high.

Second, when a firm's owners do not effectively control the incumbent managers, the managers may seek to preserve their autonomy or their jobs, by substantially raising the costs of changing the firm's form of ownership. The managers are particularly likely to be successful in this regard where, as in many cooperative and mutual firms, shares in ownership are not freely marketable.

Both of these factors produce inertia in the selection of organizational forms. This inertia is more pronounced for some forms of ownership than others. As we shall see, there are industries in which anachronistic forms of ownership have remained firmly embedded long after they have lost their original efficiency advantage over other forms.

Applying the Calculus

Although the particular categories of costs described here do not exhaust all the efficiency considerations relevant to ownership, they usefully organize those that appear most important. Ignored here are some other considerations, such as the "horizon problem," the problem of "perverse supply response," and the tendency of cooperatives to "degenerate" into investor-owned firms, that have sometimes been emphasized in the literature but that do not seem to play a fundamental role in determining patterns of ownership. These latter considerations will be discussed later in the context of particular industries that illustrate the issues involved.[24]

The chapters that follow show how tradeoffs among the various costs described here determine the structure of ownership in particular industries. In anticipation of those analyses, some general comments about these tradeoffs are in order.

As noted in Chapter 1, the efficient assignment of ownership minimizes the sum, over all the patrons of the firm, of the costs of market contracting and the costs of ownership. If the class of patrons for whom the costs of market contracting are highest is also the class for whom the costs of ownership are lowest, then those patrons are unambiguously the most efficient owners. This is often the case for small businesses.

Farms in the staple grain crops, such as wheat and corn, are obvious examples. It is not costly to borrow most of a farm's capital on the market, because the land, equipment, and crops can be pledged as security. Nor is it costly to sell the farm's products on the market, since they are simple, standardized, and easily evaluated by their purchasers (and since, to the extent that the purchasers have market power, this can be dealt with by farm-owned marketing cooperatives). Most farm inputs are also sufficiently simple and standardized to permit their purchase on the market with little cost, and farm-owned supply cooperatives provide a good solution where this is not the case. In contrast, hiring all of the labor for the farm on the market would generally lead to serious inefficiency owing to the difficulty of monitoring farm work—essentially a problem of asymmetric information—and this problem cannot be solved by having the farm own its workers. These costs of labor contracting can, however, largely be avoided by giving ownership of the farm to the family that provides most of the farm's labor. As for the costs of ownership, two of the three principal categories of those costs—the costs of monitoring managers and the costs of collective decision making—are obviously low for family farms. The chief cost of family ownership is risk bearing, and this can be mitigated by passing risk on to the market (via futures contracts), to insurers (via crop insurance), to the government (via price supports), and to creditors (via default).

Yet frequently—and especially in large-scale enterprise where the relevant classes of patrons are sizable—the efficient assignment of ownership is not so obvious. One reason is that, when the costs of market contracting are high for a given class of patrons, the costs of ownership are often high too, and for much the same reason: because it is costly for the patrons in question to become informed about how well the firm is serving them. Life insurance policyholders in the early nineteenth century provide an example we shall return to. Contracts alone were insufficient to assure the policyholders that their insurance company would ultimately pay off on their policy, yet the policyholders were too numerous and dispersed to exercise meaningful control over their insurance company if they owned it collectively.

Such patrons are often efficient owners, despite their high costs of ownership. Even if they cannot monitor the firm's management effectively, and thus cannot exercise much control over the firm beyond that available simply through market transactions with the firm, it does not follow that there is no substantial gain from having those patrons own the firm. To use Albert Hirschman's felicitous terminology,[25] it can be efficient to assign ownership to a given class of patrons even if, for those patrons, voice adds little to exit in controlling the firm. An important reason for this is that, by virtue of their ownership, the patrons are assured that there is no other group of owners to whom management is responsive. It is one thing to transact with a firm whose managers are nominally your agents but are not much subject to your control; it is another to transact with a firm whose managers are actively serving owners who have an interest clearly adverse to yours.[26]

In short, the costs of contracting for a class of patrons may be substantially reduced by making those patrons the owners even if they will only be very passive owners. Thus life insurance companies in the early nineteenth century were typically owned by their policyholders. Large U.S. industrial corporations in the twentieth century are arguably another example, as will shortly be discussed.

In the extreme, when both the costs of market contracting and the costs of ownership are exceptionally high for a given class of patrons, the efficient solution is sometimes to assign ownership to none of the firm's patrons but instead to form an unowned, or nonprofit, firm. Making owners of anyone other than those high-cost patrons would inefficiently threaten those patrons' interests. Yet making those patrons owners would result in no meaningful reduction in the agency costs of delegated management, while leading to useless administrative burdens (such as keeping track of and communicating with the nominal owners) and running the risk that the members of some subgroup

will succeed in using their authority as owners to disadvantage fellow patron-owners who are less well positioned.

In any event, as we shall see in Chapters 13–15, the distinction between nonprofit firms and firms owned by patrons who are very poor monitors is often negligible. Indeed, the tenuous character of that distinction is an important theme even in the following chapter on investor ownership.