

Governing the Firm

Workers' Control in Theory and Practice

GREGORY K. DOW

Simon Fraser University, Canada



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Gregory K. Dow 2003

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2003

Printed in the United Kingdom at the University Press, Cambridge

Typeface Times Ten 10/13 pt. *System* L^AT_EX 2_ε [TB]

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication Data

Dow, Gregory K., 1954–

Governing the firm : workers' control in theory and practice / Gregory K. Dow.
p. cm.

Includes bibliographical references and index.

ISBN 0-521-81853-2 – ISBN 0-521-52221-8 (pb.)

1. Industrial management – Employee participation. 2. Employee ownership.
3. Corporate governance. 4. Industrial organization (Economic theory) I. Title.
HD5650 .D64 2002

338.6 – dc21 2002023440

ISBN 0 521 81853 2 hardback

ISBN 0 521 52221 8 paperback

Contents

<i>List of Tables and Figures</i>	<i>page x</i>
<i>Preface</i>	xi
1 Introduction	1
1.1 Economic Systems	1
1.2 The Control Dimension	4
1.3 Looking for Clues	8
1.4 A Projected Synthesis	12
1.5 The Plan of the Book	18
2 Normative Perspectives	23
2.1 Why Care About Workers' Control?	23
2.2 Equality	24
2.3 Democracy	27
2.4 Property	32
2.5 Dignity	34
2.6 Community	38
2.7 The Author Shows His Cards	41
3 Workers' Control in Action (I)	45
3.1 Surveying the Terrain	45
3.2 The Plywood Cooperatives	50
3.3 The Mondragon Cooperatives	57
4 Workers' Control in Action (II)	67
4.1 The Lega Cooperatives	67
4.2 Employee Stock Ownership Plans	76
4.3 Codetermination	83
5 Conceptual Foundations	92
5.1 The Theory of the Firm	92

5.2 The Nature of Authority	98
5.3 The Locus of Control	101
5.4 Why Firms Cannot be Owned	107
5.5 Asset Ownership	110
5.6 Residual Claims	114
6 Explanatory Strategies	117
6.1 The Symmetry Principle	117
6.2 The Replication Principle	119
6.3 Transaction Costs	121
6.4 Optimal Contracting	126
6.5 Adverse Selection	129
6.6 Repeated Games	132
6.7 Historical Explanations	135
6.8 Cultural Explanations	138
6.9 The Strategy to be Pursued	140
7 A Question of Objectives	142
7.1 What do LMFs Maximize?	142
7.2 The Illyrian Firm	143
7.3 Membership Markets and Labor Markets	148
7.4 Membership Markets and Stock Markets	152
7.5 Imperfect Membership Markets	156
7.6 What does the Evidence Say?	161
7.7 Some Lessons	163
8 Views from Economic Theory (I)	165
8.1 Explaining the Rarity of Workers' Control	165
8.2 Asset Ownership: Incentives and Information	166
8.3 Asset Ownership: Bargaining and Investment	168
8.4 Can Asset Specificity Explain the Rarity of Workers' Control?	172
8.5 Work Incentives without Risk Aversion	173
8.6 Work Incentives with Risk Aversion	178
8.7 Can Work Incentives Explain the Rarity of Workers' Control?	179
9 Views from Economic Theory (II)	185
9.1 Capital Constraints	185
9.2 Debt Financing	186
9.3 Equity Financing	188
9.4 Can Capital Constraints Explain the Rarity of Workers' Control?	189
9.5 Portfolio Diversification	193
9.6 Can Portfolio Diversification Explain the Rarity of Workers' Control?	197

9.7 Collective Choice	200
9.8 Can Collective Choice Explain the Rarity of Workers' Control?	204
10 Transitions and Clusters	207
10.1 Organizational Demography	207
10.2 Formation Rates	207
10.3 Worker Takeovers	212
10.4 Degeneration	221
10.5 Investor Takeovers	224
10.6 Survival Rates	226
10.7 Business Cycles	228
10.8 Clusters	230
11 Toward a Synthesis	234
11.1 The Causal Tapestry	234
11.2 Credible Commitment toward Labor	240
11.3 Credible Commitment toward Capital	245
11.4 The Composition of Control Groups	249
11.5 The Commodification of Control Rights	251
11.6 Intellectual History and Current Debates	253
11.7 Is Workers' Control a Unitary Phenomenon?	256
12 Getting There from Here	260
12.1 Practical Considerations	260
12.2 A Modest Proposal	262
12.3 Reassuring Shareholders	273
12.4 Governing Firms	276
12.5 Trading Jobs	279
12.6 Sample Calculations	282
12.7 The Long and Winding Road	287
<i>References</i>	291
<i>Index</i>	313

List of Tables and Figures

Tables

1.1 Ownership and control as independent dimensions of firm organization	<i>page 3</i>
5.1 A classification of governance structures: Unbundled input supply	111
5.2 A classification of governance structures: Bundled input supply	113
12.1 Years needed to achieve workers' control with dividends reinvested in purchase of equity shares	284
12.2 Years needed to achieve workers' control with dividends paid out as cash	285
12.3 Steady-state employee share ownership with dividends reinvested and equity claims sold to outsiders on departure	286

Figures

7.1 Income maximization in the Illyrian firm	145
--	-----

Introduction

1.1 Economic Systems

People care about the organization of firms. Their concerns include alienation in the workplace and participation in decision-making, wage and job security, the risks associated with employee stock ownership and profit-sharing, and a host of other matters. Alienated workers are unlikely to contribute their best efforts to the success of the firm. Employees may wonder whether the information they reveal now will be used against them later, or whether authority delegated today will be revoked tomorrow. Information technology may make worker knowledge the most important input to the production process, but workers who fear downsizing may be reluctant to invest in skills that would only be valuable at their current job. Globalization expands export markets but simultaneously poses risks to wages and jobs. Employee stock ownership is often claimed to increase productivity, but employees may well hesitate to place their life savings in the hands of their own employer.

Many of these dilemmas would be mitigated or eliminated if workers had ultimate control over the firms to which they supply their labor. Presumably if workers ran their own firms they would feel less alienated; they would willingly disclose information that might improve efficiency; they would be less worried about layoffs, downsizings, or plant closures; and they would closely monitor how their savings were used by the firm. Why not, then, adopt an economic system in which workers rather than investors hold ultimate authority within firms? Why not implement a system of workers' control?

A prudent worker, investor, or citizen would no doubt ask a few questions before signing on to a program of this kind. Are any firms controlled by workers today? What is their track record? How are they organized? And perhaps the most important question: If workers' control is so advantageous, why is it so rare? This book seeks to answer these questions, especially the last one.

It may seem surprising that economists have not yet come to any consensus on the reasons for the rarity of workers' control. Despite much attention within the profession to the organization of firms, the question of why large firms are conventionally controlled by investors rather than workers has not been high on the economic research agenda, perhaps for the same reason that fish do not study water. It is also surprising that economists have not been more interested in policy questions relating to workers' control: for instance, the subsidies widely granted for employee stock ownership, or the support local governments sometimes provide to employee buyouts of failing firms. By and large, the profession has been content to view such measures as ad hoc political exercises that do not require any novel theoretical or conceptual framework.

I will return to the economic literature on workers' control later in this chapter, but first a short historical digression may be useful in setting the stage. The twentieth century witnessed a global struggle between two alternative systems: capitalism and socialism. As normally defined, capitalism and socialism differ along the dimension of asset ownership. Under capitalism, the physical assets used in production are owned by private investors, while under socialism they are owned by the state. Soviet central planning is, of course, discredited (Kornai, 2000), but socialism more broadly remains identified with the public ownership of capital goods (Stiglitz, 1994). Those in the West who favor this alternative have largely embraced some form of market socialism: that is, a combination of public asset ownership with reliance on market forces to handle most resource allocation problems (Le Grand and Estrin, 1989; Bardhan and Roemer, 1993; Roemer, 1994).

Economic systems, however, vary along another dimension that is quite distinct from the private or public ownership of physical assets: control over production activities. Here the distinction is between firms controlled by capital suppliers and those controlled by labor suppliers. The claim that workers should control firms democratically differs sharply from the claim that productive assets should be owned by the state. Nevertheless, writers on the political left have sometimes advocated both ideas simultaneously, perhaps believing that if one alternative to capitalism is good, then two

Table 1.1 *Ownership and control as independent dimensions of firm organization*

		Asset ownership	
		Private	Public
Control rights	Control by capital	Capitalist firm	Socialist firm
	Control by labor	Laborist firm	Self-managed firm

must be better. An unfortunate result is that the tensions between these two modes of organization have often been glossed over.

Table 1.1 shows that four separate economic systems can be distinguished by treating the ownership of physical assets and control over production as independent dimensions of firm organization. The terms “capitalist firm” and “socialist firm” in Table 1.1 are standard. Although differing in the dimension of asset ownership, both assign ultimate managerial control over the firm to capital suppliers, located either in the private sector or in the state. The “laborist” firm is a coinage meant to suggest that assets are privately owned but firms are controlled by their workforces rather than investors. The fourth cell can be identified with the Yugoslav system of workers’ “self-management” that developed during 1949–91, where the assets of firms were viewed as social property (in practice, owned by the state) but workers within each enterprise had substantial control over their use.

Like public ownership, the idea of workers’ control has a long pedigree, motivating a diverse assortment of organizational experiments ever since the Industrial Revolution. By contrast with public ownership, however, there have been few attempts to impose workers’ control on entire societies by political fiat. The only major historical example along these lines is that of the former Yugoslavia. It is difficult to draw lessons from this experiment because the effects of workers’ control were confounded with those of price regulation, poorly defined property rights, and political authoritarianism. It is therefore impossible to use the Yugoslav experience as a basis for predictions about the performance of workers’ control in a setting of free markets, private ownership, and political democracy. Readers interested in the Yugoslav case may want to examine Horvat (1971), Estrin (1983), Sacks (1983), and Stephen (1984: ch. 5). A concise overview of the entire 1949–91 period is provided by George (1993: ch. 3).

My working definition of capitalism is that it is an economic system in which most large enterprises are under the ultimate control of

capitalists – that is, individuals or groups who supply capital to firms. The relevant alternative for this book is “laborism,” in which most large firms are under the ultimate control of labor suppliers. In each case, private asset ownership is taken for granted, so I will be concerned only with the left column of Table 1.1. The term “socialism” will be discarded because of its association with public ownership, which is irrelevant to the question at hand.

It is vital to stress at the outset that private ownership of productive assets does not necessarily imply capitalism in the sense that firms are actively managed by the suppliers of capital (or their agents). In principle, worker-controlled firms can lease physical assets from their owners on a competitive market without ceding control over the production process, just as conventional firms routinely lease computers, offices, aircraft, machinery, and land. There is nothing about private asset ownership that implies routine control over the ways in which assets are used, any more than owning an apartment building implies that a landlord can supervise the daily lives of tenants. There are of course practical reasons why various assets are seldom leased, and these will be explored later. But in theory, workers could hire an asset from its owner, rather than being hired by the asset owner. Therefore, ownership does not imply control.

If it is not feasible to lease physical assets, a worker-controlled firm could finance collective asset ownership using debt contracts under which workers retained control over the firm’s activities except in extreme circumstances such as default on contractual obligations. Workers could also draw upon their personal savings to finance collective asset ownership by the firm, in which case they would become suppliers of both capital and labor. Such a firm qualifies as “worker-controlled” or “labor-managed” as long as votes are proportional to labor supply rather than capital supply.

The fact that capitalists frequently do exercise control over the activities of workers is not a logical necessity but an empirical fact requiring an explanation. A system of private asset ownership and decentralized markets need not have led to capitalist firms. However, despite occasional successes, worker-controlled firms occupy only a marginal position in the economies of North America and Western Europe. This presents a puzzle: Why does workers’ control remain rare? Or, in another formulation: Why does capitalism persist?

1.2 The Control Dimension

Modern economic theory suggests that a firm should be seen as a set of incomplete contracts among input suppliers (Williamson, 1985; Grossman

and Hart, 1986; Milgrom and Roberts, 1990a). In this view, the right to make decisions not previously determined by contracts must be assigned to some person or group. Such control rights could be used to determine a firm's product line, investment strategy, wages and employment, production methods, or working conditions, for example. In most large enterprises, these decisions are made by managers who are ultimately hired and fired by a board of directors. The directors in turn are chosen by shareholders.

This is the archetypal capitalist firm. In its "laborist" counterpart, managers may also be hired and fired by a board of directors, but the board is chosen by, and is accountable to, the workforce. Firms of both types must somehow induce managers to pursue the goals promulgated by the board of directors, a problem that will become relevant on occasion in later chapters. But the key question here is why ultimate control over the firm normally rests with investors (or their representatives) rather than with workers (or their representatives).

Because a variety of ambiguities can arise in practice, including the possibility that workers may supply both labor and capital to some firms, it is necessary to refine the rough descriptions given above. A capital-managed firm (KMF) is defined as an enterprise in which ultimate control is allocated by virtue of, and in proportion to, capital supply, while a labor-managed firm (LMF) assigns control by virtue of, and in proportion to, labor supply. Other complications – for example, involving codetermination and collective bargaining – are addressed in Chapter 5, but these preliminary definitions will suffice for the moment.

The reference to "ultimate" control is intended to bypass the problem of managerial incentives. Even if the shareholders can agree among themselves on the objectives of the firm, their control over the board of directors is limited, and directors likewise have limited control over top managers. This is the famous "separation of ownership and control" first highlighted by Berle and Means (1932), which is not to be confused with the very different distinction between ownership and control drawn in Table 1.1. Despite these incentive problems, however, it is nevertheless significant that the board of directors can dismiss the top management of the firm and that shareholders can replace the directors.

It could be objected that a focus on ultimate control is misplaced because workers can influence the decisions of their organizational superiors even in capitalist firms. For example, employees might exercise influence through committees dealing with innovation, product quality, working conditions, or grievances. More covertly, they might exert social pressure on managers, give them false information, or bribe them (Milgrom and

Roberts, 1990a). It may therefore be misleading to say that such firms are controlled by investors.

I ignore this objection for the following reasons. In an investor-controlled firm, the authority delegated to employee teams or committees can always be revoked. This is not to say that such participative processes are a fraud. They play an important role in conveying information to managers, motivating workers, accelerating responses to market conditions, and taking some of the load off vertical communication channels. But to grasp the deeper structure of the firm, it is illuminating to ask who has ultimate authority, meaning authority that cannot be revoked by anyone else. It may be quite costly or difficult in the short run to revoke delegated authority – and to this extent workers do have some *de facto* control – but in the long run, investors generally have the upper hand. In capitalist firms, the shareholders directly or indirectly determine whether other individuals, such as managers and shop floor employees, will continue their association with the firm. Employees often have a good deal of informal influence, but at the end of the day they cannot fire their bosses.

Another potential objection is to deny that the large corporation is a capitalist firm, or capital-managed, by arguing that most shareholders are not really capital suppliers. The point is that when one person sells existing shares to a second person, the firm itself does not gain access to any new capital, and so (it may be argued) the new shareholder does not supply any. This argument is mistaken because the current shareholders as a group can liquidate the firm's assets, though more than a simple majority vote is generally required. If they do so, they will share in whatever funds are left after various creditors (bondholders, suppliers, tax authorities) have been paid. By refraining from liquidation, the shareholders keep part of their wealth tied up in the firm and available for its use. As individual shares are traded, each shareholder inherits the voting rights of the previous shareholder, gains an identical claim on dividends and on the firm's non-human assets in the event of liquidation, and thus occupies a structurally identical position within the firm. It therefore makes sense to say that one capital supplier replaces another.

More generally, any firm whose complexity extends beyond sole proprietorship has some governance structure that defines the organizational roles assigned to suppliers of particular inputs. Whenever there is turnover among input suppliers, the rights and duties associated with such roles must be transferred from one person to another according to the procedures specified by the governance structure. In professional partnerships, all partners must usually agree before a partner can be replaced.

Some worker-controlled firms permit members to sell their positions to replacement workers with the approval of other members. Shareholders in a publicly traded firm, however, are uniquely free to sell their roles in the firm without first obtaining permission from any other parties. A key theoretical question is why it appears to be easy to create markets for the position of “capital supplier” but more difficult to create parallel markets for the position of “labor supplier.”

Where do control rights come from? They may arise in two general ways: through contracts, laws, or constitutions on the one hand; or through social customs, conventions, or norms on the other. The first case is the most familiar. In modern societies, a group of input suppliers can typically draw up a partnership agreement or a corporate charter that says who will have the authority to make managerial decisions. This agreement or charter becomes the ultimate source of authority within the firm, and the resulting authority may be delegated to other parties such as managers. Sometimes a specific assignment of control rights to capital or labor suppliers is fixed by law (as in German codetermination, discussed in Section 4.3) or by a constitution (as in the former Yugoslavia), but there is usually considerable scope for voluntary contracting. In this setting, whether property rights over non-human assets lead to control rights over the firm as a whole depends on the nature of the contract.

A less standard way of thinking about control rights, but one that should not be dismissed *a priori*, involves social custom or non-cooperative equilibrium in the absence of a supporting legal framework. As a thought experiment, imagine that there are a number of locations at which production can occur, and that the first person to reach a production site can exclude others who arrive later, perhaps because of social convention or through physical force. If others are admitted to the site, they cannot subsequently be excluded against their will. In this environment, the first arriver might charge a price for admission, with anyone paying the price being allowed in. After a group of input suppliers has been assembled in this way, it is necessary to make coordinated decisions about production activities. At this stage, each person could threaten to leave, taking their resources with them, unless the joint decisions are satisfactory. This would give each some bargaining power over the group’s decisions, and thus a measure of control. Other input suppliers who remained vulnerable to expulsion and were easily replaced would of course lack parallel influence over production activities. If the people with bargaining power all happen to be capital suppliers, it would be reasonable to speak of a capital-managed firm, and similarly if they are all labor suppliers.

The point of this parable is conceptual. Control rights need not derive solely from legally recognized property rights or contractual agreements. I will not be concerned here with firm organization in an anthropologist's tribe or an anarchist's utopia, but clearly one could discuss control rights over firms in such situations without lapsing into incoherence. More significantly, one should not underestimate the role of social custom or convention as a widely accepted source of authority, even in contemporary societies. Contracts exist, but so do stable patterns of behavior when the courts are not watching. It often makes sense to accept the boss's authority simply because it is correctly expected that everyone else will do the same. Moreover, a boss often has ways of enforcing compliance without resorting to a lawsuit: for instance, by threatening an employee with dismissal.

The term "control rights" as used in this book does not necessarily refer to legal or contractual rights, and never refers to moral rights. The meaning is causal, not normative. References to "control rights" can always be translated as "control abilities, or capacities, or powers." To say that a group of people has control rights in a given firm means (roughly) that they can directly or indirectly allocate the resources of input suppliers located inside the firm, distribute rewards or penalties among them, and perhaps terminate their relationships with the firm. These powers often have a formal contractual basis, but this is not essential. I will return to these issues at greater length in Chapter 5.

1.3 Looking for Clues

A naive economics student might expect the discipline to have developed a cogent and empirically supported explanation for the conventional assignment of control rights to investors rather than workers. After all, this problem is one of the most basic that could be addressed by economists, with implications for labor economics, comparative economics, industrial organization, finance, and economic history, among other fields. Such a student would be disappointed. Remarkably, economists still lack a commonly accepted rationale for the prevalence of capitalist firms in market economies.

As will become painfully clear, hypotheses about the rarity of workers' control are in excess supply. A small band of empirical researchers has labored to redress the balance, but abstract modeling has outpaced the evidence. Much theoretical discussion leans toward casual storytelling rather than thoughtful analysis informed by factual knowledge. Writers

occasionally announce that they have discovered why LMFs are rare without showing any awareness that other hypotheses exist, or that accepted evidence calls their own story into question. These difficulties are exacerbated by the ideological passions the subject tends to provoke. Some writers are convinced of the virtues of workers' control, and view its rarity as evidence for a massive market failure. Others, convinced of the market's wisdom, infer from the rarity of workers' control that it lacks all merit and barely warrants discussion.

One unfortunate consequence of this disarray is that economists are left with little to say about a wide range of current policy issues. Numerous academicians and practitioners have called for increased worker participation in the management of firms, including board representation for employees (Blair, 1995; Levine, 1995). Many firms have experimented with profit-sharing plans, employee share ownership, and more employee participation in decision making (Kruse, 1993; Blair, Kruse, and Blasi, 2000; Ben-Ner, Burns, Dow, and Putterman, 2000). A few large firms in the United States have been completely taken over by their employees (see Chapters 4 and 10), often with the aid of tax incentives or loan guarantees made available by various levels of government.

These developments raise provocative questions. Is there a process of institutional evolution underway that is leading systematically to greater workers' control within the modern firm? If so, what forces are driving this process? What are the key barriers to the further spread of workers' control? Is there any rationale for policy intervention to reduce or eliminate these barriers? Is it possible to combine the advantages of investors' control in raising capital and diversifying risks with the advantages of workers' control in stimulating effort, reducing conflict, and tapping employee knowledge? Do the answers depend on the characteristics of the industry, the nature of production technology, or the size of the firm?

I do not pretend to supply complete answers to these questions, but I hope that after finishing this book, the reader will find some answers more reasonable than others. Many pages will be devoted to historical, case-study, and econometric evidence on the subject, as well as critiques of received theoretical wisdom, because I believe there is a high marginal benefit from sifting through existing data and arguments in a systematic way. Readers who seek a decisive mathematical proof or a definitive regression equation will not find it here, but they will find a gradual accumulation of evidence pointing toward specific conclusions.

Consider an example of what needs to be explained. A fact that leaps out from historical and case study accounts is that traditional

labor-managed firms are not distributed randomly across industries. They occur mainly in professional services (law, accounting), craft manufacturing (printing, leather goods, glassware, furniture), low-skill service tasks (reforestation, refuse collection, taxi service), and construction. As a general rule, LMFs do not engage in large-scale, capital-intensive production activities. There are exceptions, but this generalization is quite robust for the subset of LMFs that have been formed directly as producers' cooperatives rather than arising through conversion of a pre-existing KMF. Further facts about the distribution of LMFs across industries, conversions of KMFs into LMFs (or conversely), and the organizational designs of successful LMFs will surface later.

An ideal theory of control rights would accomplish the following things: (1) explain why LMFs account for a small share of aggregate output, assets, and employment relative to KMFs; (2) explain the observed distribution of LMFs across industries; (3) explain the organizational transitions in which a firm changes from KMF to LMF through an employee takeover, or from LMF to KMF through an investor takeover; and (4) explain the design features that successful LMFs typically have. No theoretical framework comes remotely close to satisfying all of these requirements. Instead, there is a consistent tendency among theorists to focus on one fact – the rarity of LMFs – and to emphasize one causal mechanism that allegedly explains this fact.

To see how one might begin to construct a more complete theory of the LMF from the ground up, consider momentarily a world in which there are no qualitative distinctions between inputs of capital and labor, in either physical or institutional respects. This is the level of abstraction found in microeconomic theory textbooks, for example. For reasons to be discussed in Chapters 6 and 7, in such a world there is no reason why suppliers of one input or the other would more often acquire control over firms. The distribution of KMFs and LMFs across industries would be random, or a matter of historical accident, and thus KMFs would be unlikely to dominate the economy as a whole.

Explanations for systematic differences between KMFs and LMFs must therefore be grounded on qualitative asymmetries between capital and labor as inputs. Perhaps the most fundamental asymmetry of this sort is that ownership of physical assets can be shifted from one person to another, while the capacity to supply labor services cannot be. In short, the capacity to supply labor is inalienable, both in a physical sense and also institutionally (contracts involving slavery or voluntarily accepted servitude are unenforceable in court).

This has many implications. A firm can obtain its capital inputs either as an owned stock or as a leased flow, but it can only obtain labor as a flow. A worker's time and skill cannot exceed natural bounds, but there is no upper limit on an investor's wealth. A labor supplier must often be in close proximity to other labor suppliers to join in production, and cannot be in more than one place at a time, but one person can own many physical assets in dispersed locations. Because labor inputs depend on the characteristics of the person supplying them, they are often highly heterogeneous, while financial capital is not.

Although these are suggestive observations, simply knowing that capital and labor differ, or even how they differ, is not enough. Another portion of the theory must address what the controllers of firms do. What sorts of decisions do they make? What can one say about the preferences, beliefs, and constraints of individual controllers? Do controllers face important collective action problems? How does an input supplier become a member of the controlling group or exit from this group? What does this imply about the behavior of the firm or the likely performance of KMFs and LMFs in specific environments? Juxtaposing information about the asymmetries of capital and labor inputs with information about the role of control groups within firms provides the basis for an explanatory strategy. Later I will outline the nature of this strategy in greater detail.

Nothing thus far implies that the relative economic efficiency of KMFs and LMFs must play a central explanatory role. Efficiency arguments clearly become relevant if one's theory says that surviving organizational structures will necessarily satisfy some efficiency criterion. But someone else's theory might be based on subtle forms of market failure, or game theoretic models in which multiple equilibria arise. Theories of this sort may offer explanations for the observed distribution of LMFs while asserting that everyone could be made better off through policies to encourage the creation of more LMFs. Researchers will naturally pursue diverse explanatory strategies (see Chapter 6), but perhaps all can agree that no theory should receive extra points for assuming the efficiency of market outcomes, or for assuming the contrary. A theory should be awarded a great many points, however, if it successfully explains the incidence, behavior, and design of LMFs.

An argument occasionally advanced against the usefulness of studying LMFs is that in a competitive economy, surviving firms must necessarily maximize profit. Firms that fail to satisfy this criterion, it is said, will be replaced by those that satisfy it, and thus if LMFs are to be viable at all, they must behave in a profit-maximizing way. However, for this to be

true, certain conditions must be met. First, markets must be complete (in the sense that all relevant features of a firm's inputs and outputs can be specified in binding contracts) and competitive (in the sense that perfect substitutes for any input or output can be obtained at an identical price). Second, and relatedly, economic profit must be the sole survival test. Because LMFs exist, and have been shown empirically to deviate in various ways from the behavior of comparable KMFs (see Chapter 7), these assumptions can safely be rejected.

Many economists see worker-controlled firms as a fringe phenomenon having little or no significance for the economy as a whole, and deserving a corresponding amount of research attention. This attitude is understandable but mistaken. Even if one lacks interest in any normative project to advance the cause of workers' control, surely it is important to explain why firm governance in private ownership economies is predominantly capitalist in form. This cannot be done without considering what it means for a firm to have a capitalist structure and identifying the main alternative ways in which firms could be organized. The contrast with worker-controlled firms is surely an obvious one. Moreover, such firms are not hypothetical entities. They exist, they have often survived in competitive environments for long periods of time, and a large amount of information is available about them. The proof of the pudding is in the eating, but I believe that a systematic attempt to understand the successes and failures of workers' control can shed light on capitalism itself, including the reasons for its broad success and its continuing limitations. As a bonus, the reader may gain an illuminating vantage point from which to ponder the economic theory of the firm.

1.4 A Projected Synthesis

I will not undertake a full synthesis of theoretical ideas about LMFs until Chapter 11, but it would be unfair, and unrealistic, to expect readers to wait that long without a glimpse of the larger picture. This section provides a preliminary sketch. Many causal linkages and qualifications are omitted at this stage for expository convenience.

Firms will be viewed as coalitions of input suppliers characterized by incomplete or missing contracts, so control over production activities must be assigned to some person or group. The only candidates for control are capital and labor suppliers. I am not concerned in this book with firms controlled by consumers or suppliers of raw materials, for example. I assume that within limits, firm controllers can enforce their decisions

by means of binding contracts or some type of self-enforcement. The latter could involve threats of retaliation in situations of non-compliance, perhaps through wage reductions or dismissal.

Firms can be distinguished from markets by the absence of short-term bargaining, among other things, so individual agents do not normally haggle or attempt to reprice their entire relationship with the firm whenever controllers issue new instructions. Obstacles to short-term bargaining include the time costs involved, private information, collective-action problems, and the inability of controllers to make credible promises about their own future behavior. Whatever the reason, the lack of such bargaining makes it difficult or impossible for the people affected by a control group's decisions to alter those decisions, although they can exit to the market if sufficiently dissatisfied. Individuals or groups who enjoy control rights can therefore impose uncompensated costs or benefits on other coalition members. One implication is that the Coase Theorem (Coase, 1960), which says that bargaining will allocate resources efficiently if transaction costs are zero, does not apply within the firm.

Controllers take account of their own interests directly, but at best will internalize the interests of other input suppliers in a partial and indirect manner. The danger associated with centralized authority in firms, as elsewhere, is that it can be abused in a self-interested way through decisions that inflict costs on other agents (Dow, 1987). If the investor-controllers in a KMF can discover a way to increase their payoffs while imposing costs on employees, they will take advantage of this opportunity. Likewise, the worker-controllers in an LMF will abuse outside investors if they find this advantageous.

Some of the most significant forms of abuse have an intertemporal dimension. It is frequently difficult for a control group to credibly reassure non-controllers that they will be rewarded later for sacrifices made today. This is not always impossible: If the parties place enough weight on future payoffs, promises of this kind may become credible (see Section 6.6). But agents do discount the future, and reputation is often an imperfect safeguard. In some settings, the temptation to abuse non-controllers may be too great to resist. Problems of intertemporal credibility are therefore a prime hunting ground in looking for behavioral asymmetries between KMFs and LMFs.

Consider a stylized fact about LMFs. When created *de novo*, rather than through an employee buyout of a KMF, they rely almost exclusively on financing from the personal savings of members and the retained earnings of the firm itself. Loans from banks are rare unless there are easily

marketed structures or inventories that can be used as collateral. Non-voting equity shares are almost never used as a way of attracting capital or transferring risk. A reasonable conjecture is that outside investors fear the consequences of putting funds into an enterprise over which they have no direct control (Williamson, 1985: ch. 12; Putterman, 1993). Other things being equal, this predicts that LMFs will be rare when (1) workers are poor; (2) the industry is capital intensive; (3) physical assets cannot easily be leased; (4) the required assets are too specialized to serve as collateral on loans; and (5) the firm is unlikely to engage in repeated transactions with investors.

KMFs may have similar trouble making credible commitments to their employees. Workers frequently make large investments whose value depends on the future behavior of their employer. Such investments include decisions to locate in a specific neighborhood or to acquire certain skills. Investments of this sort imply that workers will receive quasi-rents, which are the difference between actual wages and the worker's next best alternative on the outside labor market. These quasi-rents are at risk in a number of ways: The firm may withhold promised salary increases, it may innovate in ways that make current skills obsolete, or it may shut down its local office and shift operations elsewhere. In each case, there is a danger that the investors who control the KMF will ignore costs to workers in the form of lost wages or unemployment.

Another kind of commitment problem involves the incentives of employees to supply information. It is often observed that workers are reluctant to reveal their true abilities or the productivities of their jobs to an employer. The reason is that employers can exploit this knowledge to make incentive systems more stringent, perhaps by reducing piece rates or raising output quotas, a phenomenon called the ratchet effect (Weitzman, 1980; Dearden, Ickes and Samuelson, 1990; Allen and Lueck, 1999; Carmichael and MacLeod, 2000). Various distortions can occur as a result, including low effort levels and opposition to innovation. Alternatively, the employer may construct an incentive scheme that induces efficient effort levels but only by giving rents to productive workers, which is costly from the employer's standpoint.

The LMF has advantages in solving both problems. The worker-controllers of the firm internalize the quasi-rents they derive from specialized investments in skill or location, and take these quasi-rents into account in making shutdown decisions. The ratchet effect is more complicated, but there are reasons to believe that LMFs are less likely to suffer from such difficulties. Workers are typically well informed about the