

Why Organizations Are Such a Mess
(and What an Economist Might Do About It)

Robert Gibbons*

MIT's Sloan School and NBER

rgibbons@mit.edu

First version: September 7, 1999

This version: March 23, 2000

Rough Draft

Comments Welcome

• This essay is a rough draft of the introductory chapter for a doctoral text on organizational economics. As such, it draws unabashedly on several surveys and papers I have written or co-authored, including Gibbons (1997, 1998a, 1999, 2000) and Baker, Gibbons, and Murphy (1999). Although the text will be rooted in economics, I would like it to both draw on and reach the wide range of fields that study organizations. Accordingly, parts of this essay and the text will be rudimentary for some readers but unfamiliar or even controversial to others. Furthermore, my attempt at breadth causes me to sacrifice depth in many places. Comments are therefore eagerly solicited not only on this essay but also on the course syllabus that will be the basis for the text, which is available upon request. I apologize in advance to anyone whose work is under-emphasized or omitted in either the essay or the syllabus; all I can say is that I am all ears.

Why Organizations Are Such a Mess

(and What an Economist Might Do About It)

by

Robert Gibbons

For two hundred years, the basic economic model of a firm was a black box: labor and physical inputs went in one end; output came out the other, at minimum cost and maximum profit. Most economists paid little attention to the internal structure and functioning of firms or other organizations. During the 1980s, however, the black box began to be opened: economists (especially those in business schools) began to study incentives in organizations, often concluding that rational, self-interested organization members might well produce inefficient, informal, and institutionalized organizational behaviors.

This recent economic research on internal organization complements the more established literature on the economic theory of the firm launched by Coase (1937), who argued that firms will exist only in environments in which firms perform better than markets could. To create space for firms, Coase suggested that some environments might be plagued by “transaction costs” that cause markets to perform inefficiently. After several quiet decades, Williamson (1975) took two important steps: identifying some of the conditions that create transaction costs, such as imperfect contracts, and suggesting that firms might deal with these conditions more effectively than markets could because firms can use relational contracts, as envisioned in Simon’s (1951) theory of the employment relationship. The resulting transaction-cost economics (see also Williamson, 1985, 1996) has made substantial progress on subjects such as vertical integration, supplier relationships, and complex contracts. As these subjects suggest, however, the transaction-cost literature has focused disproportionately on activities at the boundary of the firm, paying much less attention to the firm’s internal design and functioning.

In this introductory chapter I make three claims about recent economic models of internal organization. The first is that Coase’s argument has not only much-studied implications for the boundaries of firms but also long-dormant implications for their internal functioning. If firms arise only in environments in which transaction costs would cause markets to perform imperfectly, then it is one thing to assert that firms may perform better than markets would, but quite another to assert that firms will perform perfectly.

That is, firms may be second-best (*i.e.*, the best that can be achieved) but their internal functioning seems unlikely to be first-best (*i.e.*, the best that can be imagined). After all, why should firms be oblivious to conditions that wreck markets?

Of course, organizational sociologists have long appreciated that organizations are typically not well-oiled machines. For example, the classic case studies by Blau (1955), Crozier (1964), Dalton (1959), Gouldner (1954), and Selznick (1949) depict organizations that differ radically from a hypothetical Weberian bureaucracy, with its “precision, speed, expert control, continuity, discretion, and optimal returns on input” (Merton, 1940: 561). Instead, “rules are often violated, decisions are often unimplemented, ... and evaluation and inspections systems are subverted” (Meyer and Rowan, 1977: 343). Moreover, “informal structures deviate from and constrain aspects of formal structure, and ... the organization’s intended, rational mission [is undermined] by parochial interests” (DiMaggio and Powell, 1991: 12).

My second claim is that recent economic models of internal organization fit reasonably well with this post-Weberian conception of life inside organizations. I give fleeting summaries of several such models below (and richer accounts throughout the book). Consistent with Coase’s argument, these models involve some kind of transaction cost. Some follow Williamson (1975, 1985), emphasizing imperfect contracts and specific investments, but others draw additional inspiration from strands of economic theory developed since the 1970s, such as agency theory, repeated games, and information economics (as illustrated below).

I summarize these recent economic models of internal organization by saying that they explain why organizations are “a mess but not a mystery.” One of my central goals in writing this book is to motivate and interpret this summary statement in great detail. Let me begin the process here. By “mess” I mean that the predicted organizational outcome is not first-best (*i.e.*, the outcome is worse than can be imagined). That is, I reserve the right to call an organization a mess even if it is the best that can be achieved (second-best), as well as if it is worse (third-best). For second-best organizations, an alternative title for this chapter would be “Why organizations seem so inefficient (but there isn’t anything an economist can do about it).” I use “mess” rather than “seemingly inefficient” to cement in economists’ minds the sociologists’ post-Weberian view that unimplemented decisions, subverted inspections, parochial interests, and undermined missions are ubiquitous in organizational life.

By “not a mystery” I mean two things—one micro, the other macro. The micro meaning is that in these models each person takes actions that are optimal for him or her, given the (formal and informal) incentives he or she faces. In this sense, the behavior of

each individual is not mysterious.¹ The macro meaning is that these economic models analyze environments in which any organizational design would encounter transaction costs, including the second-best organizational design that minimizes these costs. In these two senses, the design and performance of messy organizations is not mysterious.

Section 1 elaborates on my first two claims: that Coase's argument has implications for internal organization, and that recent economic models fit reasonably well with organizational sociology's conception of life in organizations. Section 1.A offers a brief and selective history of the economics of internal organization; Section 1.B sketches recent models of pay for performance, promotion rules, organizational politics, corporate culture, and herd behavior; and Section 1.C suggests one path for future research on organizations that are a mess but not a mystery. Taken as a whole, Section 1 paints a bleak picture of organizational performance. In Section 2, therefore, I turn to my third claim: that recent economic models of relational contracts have begun to suggest how organizations can improve their performance, even in environments that cause both markets and firms to perform imperfectly. This claim begins to explain how an economist might move an organization from the third- to the second-best.

As Barnard (1938), Simon (1947), and many others have noted, firms are riddled with relational contracts—informal agreements that powerfully affect the behaviors of individuals within the firm. Virtually every collegial and hierarchical relationship in organizations involves important relational contracts, including informal *quid pro quos* between co-workers and unwritten understandings between bosses and subordinates about task-assignment, promotion, and termination decisions.² Even ostensibly formal processes such as compensation, transfer pricing, internal auditing, and capital budgeting

¹ Some readers (especially non-economists) may harbor the misconception that economic models predict that rational, self-interested people will achieve efficient outcomes. In fact, an economic model's prediction of efficiency rests as much on its assumptions about the environment as on those about the people. In a social dilemma or commons problem, for example, each person's incentive is to free-ride (*i.e.*, to contribute only as much as is warranted by the resulting increase in his or her own benefit, ignoring the benefits to others), so rational, self-interested individuals are predicted to achieve an inefficient group outcome. The recent economic models of internal organization are like those of social dilemmas: rational, self-interested people are again predicted to produce an inefficient outcome (or "mess").

² For early commentary along these lines, see the classic case studies cited above—by Blau, Dalton, Gouldner, and Selznick—that inspired American sociology's departure from Weber's emphasis on formal organizational structures and processes. Granovetter (1985: 502) offers a more recent assessment: "The distinction between the 'formal' and the 'informal' organization of the firm is one of the oldest in the literature, and it hardly needs repeating that observers who assume firms to be structured in fact by the official organization chart are sociological babes in the woods."

often cannot be understood without consideration of their associated informal agreements.³

But business relationships are also riddled with relational contracts. Many transactions do not occur in a pure spot market between buyers and sellers who pass (goods) in the night. Instead, supply chains often involve long-run, hand-in-glove supplier relationships through which the parties reach accommodations when unforeseen or uncontracted-for events occur.⁴ Similar relationships also exist horizontally, as in the networks of firms in the fashion industry or the diamond trade.⁵ Whether vertical or horizontal, these relational contracts influence the behaviors of parties across firm boundaries.

Simply put, close observers have long argued that relational contracts are crucial for conducting many transactions, whether within or between firms. Section 2 discusses the modeling, applications, and future of this idea in economics. Section 2.A introduces the basic repeated-game model that economists now use to analyze relational contracts; Section 2.B discusses how this model can be used to revisit the Coase-Williamson comparison of markets versus firms, taking into account the ubiquity and importance of relational contracts in both domains; and Section 2.C suggests how future elaborations of this model may help economists grapple with unfamiliar topics such as management and leadership.

To conclude this introductory chapter, Section 3.A outlines the rest of the book, showing which parts expand on the models summarized in Sections 1.B and 2.B of this chapter; Section 3.B begins what I hope will be an ongoing discussion concerning the limitations of a strictly economic approach to organizations; and Section 3.C concludes by discussing the value and prospects of this economic approach.

³ See Lawler (1971) on compensation, Eccles (1985) on transfer pricing, Dalton (1959) on internal auditing, and Bower (1970) on capital budgeting. See also Blumenstein and Stern (1996) on how the 1700-page contract between General Motors and the United Auto Workers has important gaps that are covered by informal agreements.

⁴ Macaulay (1963) was one of the first to emphasize the importance of such “non-contractual relations” in various businesses. See also Dore’s (1983) classic account of relationships between firms in the Japanese textile industry. Kogut, Shan, and Walker (1992) suggest the prominence of such relationships by relabeling the make-or-buy decision as “The Make-or-Cooperate Decision.” Eccles (1981) describes “quasifirms” in the construction industry—long-run relationships between general contractors and independent, specialized subcontractors. Other examples permeate the literature.

⁵ In “Neither Market Nor Hierarchy: Network Forms of Organization,” Powell (1990) describes a variety of other examples and emphasizes their differences from spot markets and firms. See Podolny and Page (1998) for a summary and critique of the growing literature on networks.

1. Why Organizations Are Such a Mess

If Coase is correct (that firms will exist only where they perform better than markets could) then it is a short step to argue that inefficiency inside organizations will be typical, not exceptional. Figure 1 illustrates such an argument, by plotting the declining social surplus generated by market governance as transaction difficulty increases (*e.g.*, as imperfect contracts and specific assets become more problematic) and also the declining social surplus generated by firm governance as transaction difficulty increases.⁶ At the critical value of transaction difficulty indicated by the dotted line, markets and firms are equally effective governance structures.




Figure 1. Coase (1937) Meets Heckman (1976).

I intend this figure to be only suggestive. That is, the figure is an informal attempt to depict Coase's argument: transactions to the right of the dotted line will be governed by firms, those to the left by markets. But if this depiction is even roughly right then it follows that inefficient internal organization is typical, not exceptional: compare the observed effectiveness of firms (to the right of the dotted line) with the observed effectiveness of markets (to the left of the dotted line); the latter is superior, especially as

⁶ If we define transaction difficulty to be the collection of features that cause the effectiveness of market governance to decline then it is definitional that the curve labeled "market" in the figure declines as transaction difficulty increases. But it does not follow that the curve labeled "firm" also declines as this definition of transaction difficulty increases. I do not assert that the same collection of features always causes the effectiveness of both market and firm governance to decline. Much of the rest of the book can be interpreted as trying to analyze whether and where a figure like this makes sense.

transaction difficulty falls to zero, at which point market governance produces the efficient outcome familiar from neoclassical economics. In brief, this figure suggests that firms may live in tough environments and so be a mess (*i.e.*, seemingly inefficient but in fact second-best). This is a sample-selection argument, akin to Heckman (1976).⁷

My second claim is that economics has recently begun to produce formal models of internal organization that fit the evidence reasonably well. In these models, (second-best) inefficiency is typical, not exceptional. In this introductory chapter, I can only begin to document this claim; more theory and evidence appear in later chapters. I begin in Section 1.A with a brief and selective history of the economics of internal organization. Then, in Section 1.B, I sketch recent models of pay for performance, promotion rules, organizational politics, corporate culture, and herd behavior. All of these models illustrate “transaction difficulties” of the kind summarized on the horizontal axis of Figure 1, and consequently also illustrate imperfect transaction effectiveness of the kind summarized on the vertical axis. Together, these models also match the spirit (if not yet the details) of the evidence: the models predict inefficient, informal, and institutionalized organizational outcomes (in senses I will define below). Finally, in Section 1.C, I discuss how future economic models in this spirit might revitalize the once-active sociological literature on power and politics in organizations.

1.A A Brief and Selective History

Because the black-box model of the firm dominated economics for two hundred years, the history of economic models of internal organization is short. I find it useful to begin in 1972, with two landmark publications: Arrow and Hahn’s *General Competitive Analysis* and Marschak and Radner’s *Economic Theory of Teams*. Both books began with elegant summaries of existing work and then progressed to the authors’ new results.

Arrow and Hahn summarized two centuries of progress on the single largest research agenda in the history of economics. The central question had been posed by Adam Smith: could the relentless pursuit of self-interest by hordes of tiny firms and consumers yield anything but chaos? Shockingly, the answer was that (under certain assumptions, including perfect contracting and perfect competition—both of which will be abandoned below) not only would there be an equilibrium rather than chaos, but it

⁷ Note well that each different value of transaction difficulty on the horizontal axis represents a different transaction. The fact that one transaction has nearly zero transaction difficulty certainly does not imply that all transactions do. A distribution of transactions at various transaction difficulties could thus imply an economy with some transactions conducted within firms and others between (*i.e.*, in markets).

would be impossible to rearrange the allocation of resources in such an equilibrium to make all participants better off.

Marschak and Radner, in contrast, took a relatively new subject—the axiomatic single-person decision theory of von Neumann and Morgenstern (1944) and Savage (1954)—and applied it to a brand new area of economics. Before Marschak and Radner, the economics literature contained extremely little formal modeling of the internal organization and operation of firms. (For example, the firms that appeared in the general-equilibrium models summarized and advanced by Arrow and Hahn were black boxes.) Marschak and Radner applied decision theory in “team” settings, where different agents have different information and control different actions but all agents share a common goal (such as maximization of the firm’s profit).

Ironically, 1972 also saw the publication of the antithesis of team theory: “A Garbage Can Model of Organizational Choice,” by Cohen, March, and Olsen. Whereas team theory envisions an organization whose members compute and execute optimal communication rules to achieve efficient decisions, a garbage can is “organized anarchy.” Garbage cans are “collections of choices looking for problems, issues and feelings looking for decision situations in which they might be aired, solutions looking for issues to which they might be the answer, and decision makers looking for work” (p.1). I find the garbage-can conception of organizations to be usefully provocative: it may overstate the level of anarchy in most real organizations, but it provides desperately needed contrast to the Weberian models from team theory. For purposes of this brief and selective history, the crucial lesson from the garbage-can model (and from many of March’s other contributions) is that it is often not useful to think of an organization as a single, unified, rational decision-maker—as general-equilibrium models did for 200 years, and as team theory continued to do.⁸ Put more colorfully, many organizations look more like garbage cans than like teams (but note well the enormous middle ground between these polar extremes).

Since the early ‘70s, several theoretical developments have allowed economic models to move away from team theory, towards garbage cans. For expositional purposes, I parse these theoretical developments into the five categories given above: imperfect

⁸ As but one example of March’s contributions along these lines, consider “Information in Organizations as Signal and Symbol” by Feldman and March (1981). I read this paper as a critique of team theory, *not* of economics in its entirety. Most of the organizational behaviors Feldman and March catalogue are inconsistent with formal theories of rational choice by *single* individuals—and so are inconsistent with viewing the organization as a single, unified, rational decision-maker—but are at least partially consistent with (say) simple game-theoretic models of signaling or free-riding. Indeed, Feldman and March cite some of the classic work in information economics, and presumably would have cited more explicitly organizational models had the organizational economics literature been more developed in 1981.

contracts, specific investments, agency theory, repeated games, and information economics. The roots of the large literatures on agency theory, repeated games, and information economics are not directly relevant to our discussion; we will discuss more recent and applied models from these literatures below.⁹ But it is important for our purposes to understand the early work on imperfect contracts and specific investments. These ideas were introduced to the organizations literature by Williamson (1975) and Klein, Crawford, and Alchian (1978), with Williamson making greater progress on imperfect contracts and Klein, Crawford, and Alchian on specific investments. Williamson (1979) then presented an early synthesis of these two ideas, and christened the ensuing research stream “transaction-cost economics.”

Transaction-cost economics (TCE) has had such an impact on organizational economics that some readers may wonder whether there is any difference between TCE and the literature I summarize in this book. The two are certainly related (even intertwined), so to conclude this history, let me explain the difference I see. Beginning with Coase, TCE has primarily emphasized inefficiencies that are separate from (and perhaps logically prior to) those emphasized in this book: the inefficiencies of the market in conducting certain transactions (from which it follows that an organization might be more efficient than the market at conducting these transactions), rather than the inefficiencies of the firm in its internal design and performance. One of Williamson’s important contributions to TCE was to elucidate the roles of imperfect contracts and specific investments in such subversions of market exchange. Another was to compare the transactional efficiency of alternative governance structures, including vertical integration, non-standard contracts (*e.g.*, short-run arrangements such as “hostages”), and relational contracting. In my opinion, TCE has explored the implications of imperfect contracts and specific investments more thoroughly between firms rather than within them. What distinguishes this book from TCE, therefore, is partly a difference in methodology (through the addition of agency theory, repeated games, and information economics to the modeler’s tool kit) but more importantly a difference in focus:

⁹ By agency theory I mean models such as Holmstrom’s (1979, 1982a), in which a principal and an agent sign a contract that specifies the agent’s compensation as a function of one or more performance measures, and then the agent chooses actions in response to the incentives created by the contract. By repeated games I mean the intuitive idea that the shadow of the future can influence current behavior, both in “folk theorem” results like Fudenberg and Maskin’s (1986) for infinitely repeated games and in “reputation” models like Kreps et. al.’s (1982) analysis of cooperation in the finitely repeated prisoners’ dilemma. And by information economics I mean models involving private information, such as Akerlof’s (1970) model of used-car markets, in which a seller knows his or her car’s quality better than buyers do, Rothschild and Stiglitz’s (1976) analysis of insurance markets, in which drivers know their accident risks better than insurance companies do, and Spence’s (1973) theory of education choices, in which a worker acquires education to signal his or her productive ability to prospective employers.

transaction-cost economics has been more concerned with the boundaries of the firm than with its guts; recent economic models of internal organization reverse the emphasis.

1.B Some Recent Economic Models of Internal Organization

This section offers fleeting descriptions of recent economic models concerning five organizational issues: pay for performance, promotion rules, organizational politics, corporate culture, and herd behavior. I chose these organizational issues for two reasons. First, they illustrate the five main theoretical ideas underlying the economics of internal organization: imperfect contracts, specific investments, agency theory, repeated games, and information economics. Second, they suggest economic interpretations of three important ways that non-economist students of organizations often suggest that real organizations depart from economic models: inefficient, informal, and institutionalized organizational behaviors.

By predicting inefficient, informal, and institutionalized organizational behaviors, these models match at least the spirit of the evidence. Later chapters assess the match in more detail. But to preview the conclusion of these chapters, let me say now that I do not believe that these (or other) economic models capture all the important aspects of organizational design and performance. Nor do I believe that such economic models are the only (or best) strategy for organizational research (Gibbons, 1999: 146). Instead, my view is that such economic models come closer to capturing life in organizations than is often recognized (by non-economists who assume that economic models predict efficiency, and by economists who assume that real organizations are first-best efficient), and that such models add value in several ways to organizational research (as described in Section 3.C).

1.B.1 Pay for Performance (and Agency Theory)

In my view, the enormous literature on agency theory (now almost three decades in the making) should be construed as an attempt to progress beyond the adage “You get what you pay for.” That is, agency theory should address questions such as *what* should be paid for (*i.e.*, what performance measures should be used?) and *how* (*e.g.*, should there be objective or subjective weights on these performance measures?). From this perspective, Kerr’s (1975) classic article “On the Folly of Rewarding A, While Hoping for B” is clearly relevant.

Unfortunately, not one of the hundreds of papers on agency theory written in the 1980s could even express the insight of Kerr's title, not to mention evaluate or extend it. Indeed, I know of no evidence that the economists writing these papers knew that work like Kerr's existed in the literature on organizational behavior. Instead, the agency-theory literature explored sophisticated variations on profit-sharing, emphasizing the trade-off between incentives and insurance that results if an agent's pay is linked to the organization's profit but factors beyond the agent's control make profit uncertain.

Recent work, on the other hand, has been very much in Kerr's spirit. In particular, Holmstrom and Milgrom (1991) and Baker (1992) can be read as independently discovering problems Kerr labeled "overemphasis on highly visible behaviors" and "fascination with an 'objective' criterion" (pp. 779-80), respectively. More generally, many economists have realized that the basic agency model abstracts from important dimensions of performance evaluation that make pay-for-performance systems terribly problematic for many firms. See Gibbons (1998b) for further discussion of both the basic agency model and these recent theoretical developments.

As one of many infamous examples of difficult pay-for-performance schemes, consider the H.J. Heinz Company: division managers received bonuses only if earnings increased from the prior year; managers delivered consistent earnings growth by manipulating the timing of shipments to customers and by prepaying for services not yet received (Post and Goodpaster, 1981); such actions greatly reduced the firm's future flexibility, but the compensation system in no way addressed this issue. In keeping with the experience at Heinz, one of the central assumptions in many new models is that it is impossible to enforce a contract that makes pay (or anything else) contingent on a worker's *total* contribution to firm value; that is, these are imperfect-contract models. The rationale for this assumption is as follows: there are many ways that workers can help (or hurt) each other at a given point in time, and many ways that short-run performance can be a misleading forecast of long-run performance; as a result, the worker's total, long-run contribution to firm value typically cannot be measured precisely, especially if such measurements must be taken by a neutral outsider (say, a court) in the event of a contract dispute.

1.B.2 Promotion Rules (and Specific Investments)

Much of the work on specific investments (largely from a transaction-cost perspective) considers physical assets. Firm-specific human capital, however, is a fundamental concept in labor economics (Becker, 1964). While analyses of physical

assets often yield prescriptions about asset ownership (*e.g.*, Grossman and Hart, 1986) and hence focus on the boundaries of the firm, analyzing specific human capital naturally takes one inside the firm.

There are many ways that workers could increase their productivity with their current employer—learning more about some idiosyncratic aspect of the firm’s production, marketing, or governance, for example. Consider the subset of such investments that are inexpensive (in terms of foregone productivity while the worker is devoting time to learning): the firm would like workers to undertake these investments, and would like to promise to reward those who do so.

Like any investment, an investment in human capital is valued for its future effects—for the increases in future productivity it yields. But if it is difficult for a court to determine a worker’s *actual* contribution to firm value (as just argued in the context of pay for performance), it is even more difficult for a court to determine a worker’s *expected future* contribution to firm value. Thus, we are back to imperfect contracts: unlike Becker’s original analysis, many recent economic models assume that it is impossible to enforce a contract specifying pay or promotions based on a worker’s investments in firm-specific human capital.

We have now reached the following two-sided incentive problem: the worker is concerned that the firm cannot be trusted to reward an investment in specific human capital properly, but the firm is concerned that the worker will not invest unless such rewards are anticipated. Prendergast (1993a) analyzed whether a promotion rule might solve this problem; Kahn and Huberman (1988) considered an alternative rule, “up-or-out.”

Prendergast considered a two-job ladder with wages attached to jobs, as in some internal labor markets. (The idea is that a court could enforce a contract specifying the wage to be paid in each job, even if it could not enforce contracts contingent on specific human capital.) Each worker’s career with the firm consists of two periods: the worker begins in the low-wage job in the first period and then is either promoted to the high-wage job or kept in the low-wage job for the second period. During the first period, the worker can invest in firm-specific human capital. At the end of the first period, the firm observes the outcome of each worker’s investment and then chooses which workers to promote (if any).

The workers can rely on the firm to promote those workers who are more *profitable* for the firm when assigned to the high-wage job, but there is an important difference between being more profitable and being more productive. Promoting a worker who is slightly more productive in the high-wage job will reduce profits if the difference

in wages between the two jobs is sufficiently large. For example, if the two jobs in question are really just two job titles sharing the same underlying production technology then Prendergast's promotion rule (*i.e.*, promote those who are more profitable in the high-wage job) will not solve the two-sided incentive problem: the firm gets just as much productivity by keeping the worker in the low-wage job. Consequently, workers have no incentive to invest, even though investing would be (first-best) efficient.

Kahn and Huberman (1988) showed that an up-or-out rule may solve the two-sided incentive problem, even if the two jobs are just two job titles sharing the same underlying technology. An up-or-out rule is a promotion contract specifying that after some fixed probationary period the firm must either pay a worker a high wage or fire the worker. (Again, this contract could be enforced by a court, even if contracts contingent on specific human capital could not.) Under an up-or-out contract, the firm finds it in its interest to retain those who have made themselves sufficiently valuable (*i.e.*, more valuable than the high wage) but to fire those who have not. The workers understand that the firm will promote or fire workers in this way, and so have an incentive to invest if the high wage is high enough. For some parameter values, an up-or-out rule solves the two-sided incentive problem completely; for others, it makes partial progress (*e.g.*, it induces partial but not fully efficient investments in specific human capital).

An up-or-out rule also may have big costs, even if it solves the two-sided incentive problem. Suppose, for example, that workers who make the appropriate investment could realize any one of several different levels of specific human capital, from very low to very high. If some of the low realizations make the worker worth less than the high wage attached to the promotion then these workers will be fired. Firing these workers wastes their specific capital — a seeming inefficiency.

1.B.3 Organizational Politics (and Information Economics)

One of the classic ideas in information economics is Spence's (1973) analysis of education as a job-market signal. Signaling models have become commonplace virtually throughout economics (especially in industrial organization, labor, and macroeconomics) and in related fields such as finance, marketing, and political science; Feldman and March (1981) sketched applications of signaling to organizational behavior.

Although signaling models have received much attention in the economics literature, other aspects of information economics may be more important for understanding the structure and functioning of organizations. For example, Milgrom and Roberts (1988) introduced the idea of influence activities—attempts to manipulate

information to influence decisions to one's own benefit, even when one has no private information to signal. As an example of influence activities, consider Holmstrom's (1982b/1999) model of career concerns in labor markets: workers know that firms will use workers' outputs to draw inferences about workers' abilities, and that these inferences will in turn determine subsequent wage offers, so workers have an incentive to work hard to influence the firms' inference, even if the workers have no private information about their abilities. In Holmstrom's model, the workers' influence activities (hard work) are productive, but in many organizational contexts influence activities either distract organization members from productive tasks or merely change the distribution of organizational resources across members, without improving overall productivity.

Milgrom and Roberts suggested two ways that an organization could respond to the prospect of wasteful influence activities. First, an organization could eliminate influence activities by eliminating opportunities for influence—that is, by closing the relevant communication channels. Naturally, such a response has its costs. Second, an organization may also be able to eliminate influence activities by adjusting its internal structures and/or processes away from what would otherwise be optimal, to eliminate members' incentives to manipulate information. That is, by sufficiently distorting the organizational design, it might be possible to create a Marschak-Radner team, in which all members share a common goal. Of course, an organization could go part way down either or both of these two roads. For example, an organization could commit to limits on its discretion—perhaps by limiting the time given for debate, or by imposing other rules that partially constrain the organization's ability to respond to information provided by its members. In this case, decision makers will have the benefit of some information, but organization members will also engage in some wasteful influence activities.

Recall that Kerr first articulated issues later modeled by Holmstrom and Milgrom and by Baker. Similarly, Crozier (1964:45) lucidly described an organization that went quite far down the first of the two roads Milgrom and Roberts later proposed—towards shutting down communication entirely. In Gibbons (1999), I construct a simple model of Crozier's ideas; Section 3.C discusses the potential value of such formal modeling as a complement to detailed description and informal theorizing like Crozier's.

1.B.4 Corporate Culture (and Repeated Games)

The new work in agency theory described above emphasizes that it is often extremely difficult to measure the agent's total contribution to the firm, particularly if such measurements must be made by a neutral outsider in the event of a contract dispute.

Economists describe this difficulty by saying that the agent's contribution to firm value is not *verifiable*. Even if the agent's total contribution is not verifiable, however, it is sometimes *observable* by superiors who are well placed to understand the subtleties of the agent's behavior and opportunities. Such subjective observations of an agent's contribution to firm value may be imperfect, but they may nonetheless complement or improve on the available objective (or verifiable) performance measures.

Using models of repeated games, economists have begun to analyze relational contracts (*i.e.*, agreements enforced by reputation, rather than by the courts). The advantage of relational contracts is that they can be based on subjective observations, whereas formal contracts must be based only on objective measures (and must have prespecified weights attached to these objective measures). The disadvantage is that relational contracts must avoid creating (net) incentives for the parties to renege—they must be self-enforcing. Such repeated-game models of relational contracts begin to capture a second important aspect of life in organizations: the role of informal structures and processes in determining organizational performance.

Kreps (1990) developed a simple but influential model along these lines, using a repeated game to discuss corporate culture. In each period of the repeated game, the firm's "trading partner" (*e.g.*, a worker, a supplier, or a customer) must choose whether to "trust" the firm (*e.g.*, whether to make a specific investment). The firm then chooses either to honor or to betray this trust. If the payoffs and interest rate are such that betraying trust maximizes current profit but honoring trust maximizes the present value of current and future profit then there exists an equilibrium in which the trading partner offers trust and the firm honors it (but the trading partner would cease to offer trust if the firm ever betrayed it).

Kreps interpreted this abstract model in terms of the unexpected events that make contracts imperfect and culture instrumental. When unforeseen problems or opportunities arise, the firm's culture may help all parties decide how to respond: the culture may indicate whether this is an instance in which the firm is meant to take the long view rather than maximize short-run profit, and hence also whether this is an instance in which the trading partner ought to trust the firm. See Section 2 for quite a bit more on relational contracts in general and Kreps's model in particular.

Bull (1987) constructed a similar model to analyze subjective performance evaluation. In each period, the agent chooses how hard to work, where working harder makes high output more likely. If the agent produces high output, the firm is supposed to pay a bonus. Unfortunately, only the firm and the agent can observe the agent's output, so the firm can renege on the promised bonus if it chooses. Since the agent's output is

already in the bank, the firm decides whether to renege by weighing the current cost of the bonus against the future profit from a smooth relationship with the agent. Shocks to the value of this relationship may cause the firm to renege unexpectedly. One memorable example of the cost of renegeing on an relational contract involved a spate of departures from First Boston after Archibald Cox, Jr., having paid below-average bonuses the previous year, promised but then did not deliver bonuses on a par with those paid at comparable Wall Street firms (Stewart, 1993).

Baker, Gibbons, and Murphy (1994) combine Bull's analysis of subjective (but non-distortionary) bonuses with Baker's (1992) model of distortionary formal contracts. Using relational contracts can ameliorate the distortions caused by formal contracts; using formal contracts can reduce the size of the bonus promised in the relational contract, hence reducing the firm's incentive to renege. Such an analysis of the interplay between explicit and relational contracts is a first step towards Blau and Scott's (1962:6) observation that:

It is impossible to understand the nature of a formal organization without investigating the networks of informal relations and the unofficial norms as well as the formal hierarchy of authority and the official body of rules, since the formally instituted and the informal emerging patterns are inextricably intertwined.

In a related paper, Baker, Gibbons, and Murphy (1999) study the interplay between (formal) asset ownership and (informal) relational contracts; see Section 2.B below.

1.B.5 Herd Behavior (and Institutional Theory)

The previous four sub-sections illustrate agency theory, specific investments, information economics, and repeated games. All also involve imperfect contracts, so I have now introduced the five theoretical developments underlying the economics of internal organization. Furthermore, organizational issues such as distortionary performance measurement, up-or-out rules, and influence activities illustrate inefficient organizational behaviors, and corporate culture and relational contracts illustrate informal organizational behaviors. All that remains of the tasks for Section 1.B is to say something about institutionalized organizational behaviors.

A large literature in organizational sociology describes organizations (and other decision-makers) as taking many things for granted—adhering to “the prescriptions of myths in the institutional environment” (Meyer and Rowan, [1977] 1991: 50) rather than calculating and implementing optimal actions. I am far from ready to dismiss this idea,

but I would nonetheless like to emphasize that many of its implications are consistent with rational-choice models. For example, Banerjee (1992) analyzes an economic model of herd behavior in which rational choices produce what might be called conformity, compliance, or institutionalization; see also Bikhchandani, Hirshleifer, and Welch (1992, 1998).

Banerjee considers a sequence of decision-makers, each facing identical choices. The k^{th} decision-maker's payoff depends only on whether she makes the optimal choice, not on any aspect of the others' choices (such as how many of them make the optimal choice, or how many of them make the same choice as k). Unfortunately, each decision-maker lacks full information about which choice is optimal; instead, she begins with noisy prior information (shared by all the decision-makers) and her own private but noisy signal. The key idea is that each decision-maker observes the choices of those earlier in the sequence and so supplements her initial information with inferences about the signals held by earlier decision-makers. The resulting *optimal* decisions can exhibit herd behavior (*i.e.*, doing what others are doing rather than following one's own information).

The following example (paraphrased from Banerjee 1992: 798) shows how extreme the herd behavior might be:

There are two restaurants, A and B, that are next to each other. A publicly available restaurant guide suggests that the prior probabilities are 51 percent for A being the better restaurant and 49 percent for B. One hundred people arrive at the restaurants in sequence, observe the choices made by the people before them, and decide on one or the other of the restaurants.

In addition to knowing the prior probabilities and observing the choices made by people before them, each of these people also has a private signal (say, from a friend who has been to the restaurants) that says either that A is better or that B is better. The signal could be wrong, but is of high enough accuracy to outweigh the restaurant guide. That is, based on the prior information from the guide (which slightly favors A) and a signal that says that B is better, a person should go to restaurant B. The key question, however, will be what a person should do given (i) the prior information, (ii) a private signal in favor of B, and (iii) overwhelming evidence that other people have chosen A.

Suppose that of the 100 people, 99 have received signals that B is better but the one person whose signal favors A chooses first. Since the restaurant guide and her signal both favor A, the first person will go to A. The second person will now know that the first person had a signal that favored A (because a signal favoring B would have outweighed the slight prior advantage of A).

The second person's signal favors B, but suppose that each person's signal has the same accuracy, so a person in possession of two conflicting signals (one advising A and the other B) would ignore these two signals—they would cancel out. Thus, the second person chooses A on the basis of the restaurant guide.

Of course, the second person would have chosen A had her signal favored A (because then the guide and the first and the second signals favor A). Thus, the third person can draw no inference about the second person's signal. The third person's situation is therefore exactly the same as the second person's was: the third person has the restaurant guide and her own signal (favoring B), and can infer the first person's signal (favoring A), so she makes the same choice that the second person made (again contrary to her signal favoring B). Indeed, everyone chooses A, even though the aggregate information makes it practically certain that B is better.

Banerjee's model (and others like it; see below) should give pause to those who argue (or assume) that natural selection will produce efficiency. If one interprets the restaurant story as firms choosing technologies, for example, and supposes that firms that earn below-average profit will be selected out, then no firms go bankrupt even though no firm chooses the efficient technology.

Prendergast (1993b) develops a similar model, in which a "yes man" has an incentive to conform to the opinion of his supervisor. Extending the model to a multi-worker setting produces "group think," where individuals have an incentive to conform to the opinion of the group. A slight variation on Prendergast's one-worker model (in the spirit of Scharfstein and Stein, 1990) runs as follows.

Suppose a manager needs information from a worker (perhaps about a new production technology or marketing opportunity) but does not know the worker's ability in gathering the information (*i.e.*, a worker with higher ability is one who can gather more precise information about the underlying parameter of interest). Suppose also that the worker cannot simply pass along to the manager all the information she collects; rather, the worker can report only a point estimate of the parameter of interest. Will the worker report honestly (*i.e.*, will she report the conditional expectation of the parameter based on all the information she has gathered) or will she bias her report towards what she thinks the manager expects to hear? If the worker's future wages depend on the manager's assessment of her ability (as seems sensible and follows from several models of the labor market) then the worker will bias her report, because the manager's assessment of the worker's ability is more favorable when her report is not wildly at odds with what the

manager expects to hear. Roughly speaking, an outlandish report could be truth-telling by a very able worker but is more likely to be the result of low ability.

Now suppose there are two workers. The same ideas apply, with an extra twist: what the manager considers to be an outlandish report from one worker now depends in part on the report from the other worker. Thus, each worker now wants to conform not only to what the manager expects to hear but also to what the other worker is expected to say. This argument does not bode well for communication and decision-making in organizations.

1.B.6 Summary and Future Research

I have suggested that (a) Coase's argument has implications not only for firms' boundaries but also for their guts (namely, that organizations will be a mess but not a mystery) and (b) recent economic models of internal organization fit reasonably well with this post-Weberian conception of life inside organizations. The models summarized above were chosen to illustrate the five theoretical developments underlying the economics of internal organization, but these models also paint a bleak picture of organizational performance. In Section 2, therefore, I discuss the one way I know out: relational contracts. As described in Section 1.B.4, repeated-game models of relational contracts have been used to analyze corporate culture and subjective performance assessments. In Section 2 I go further in this direction, arguing that repeated-game models can shed new light on many structures and processes both within and between organizations, and that such models may eventually help economists grapple with such unfamiliar topics as management and leadership.

To conclude this sub-section's quick tour through recent economic models of internal organization, however, I will return to the theme of bleak organizational performance by nominating a traditional area of non-economic organizational research that I believe is ripe for economic modeling: the political perspective on organizations developed by March (1962, 1994), Pfeffer (1981), and colleagues. Consider Pfeffer's (1981: 28) comparison of the "bureaucratic" and "economic" approaches to the political approach:

In bureaucratic theories of organizations, the presumption is that through control devices such as rewards based on job performance or seniority, rules that ensure fair and standardized treatment for all, and careers within the organization, the operation of self-interest can be virtually eliminated as an influence on organizational decision making. Economic or incentive theories

of organizations argue that through the payment of wage, particularly when compensation is made contingent on performance, individuals hired into the organization come to accept the organization's goals. Political models of organizations assume that these control devices, as well as others such as socialization, are not wholly effective in producing a coherent and unified set of goals. ... *To understand organizational choices using a political model, it is necessary to understand who participates in decision making, what determines each player's stand on the issues, what determines each actor's relative power, and how the decision process arrives at a decision* [emphasis added].

Pfeffer's summary of the bureaucratic approach, in which the operation of self-interest is eliminated as an influence on organizational decision making, is reminiscent of the first economic theory of internal organization, Marschak and Radner's (1972) "theory of teams" described in Section 1.A. I certainly agree with Pfeffer that the bureaucratic/team approach misses something important (hence my assertion in Section 1.A that "many organizations look more like garbage cans than like teams"). As for the economic approach, Pfeffer's summary is accurate for its time, reflecting early work in agency theory, such as Holmstrom's (1979, 1982a). Again, I agree that there are important omissions in that approach; I much prefer the models summarized in Section 1.B.1, in which no contract could cause individuals to "come to accept the organization's goals."

Although I agree with Pfeffer's assessments of the bureaucratic and (old) economic approaches, I think the most exciting source of agreement concerns the political approach. To me, Pfeffer's summary of the political approach is fundamentally game-theoretic, if informally so. Compare his italicized statement to the following: an "extensive-form representation of a game specifies (1) the players in the game, (2a) when each player has the move, (2b) what each player can do at each of his or her opportunities to move, (2c) what each player knows at each of his or her opportunities to move, and (3) the payoff received by each player for each combination of moves that could be chosen by the players" (Gibbons, 1992: 115). In short, I read Pfeffer as a closet game-theorist!

More generally, although power and politics have almost disappeared from recent non-economic research on organizations, it seems unlikely that they have disappeared from organizations themselves, so perhaps now is the time to bring them back into the research agenda by integrating the detailed description and informal theory of the original literature with the kind of formal economic modeling summarized above. Rotemberg (1994) and Rajan and Zingales (1996) offered nice starts in this direction, but much remains to be done.

2. What An Economist Might Do About It

Problematic pay-for-performance schemes, wasted (or non-existent) investments in human capital, lobbying and other influence activities, the vagaries of subjective management practices, and herd behavior and group think all bode ill for organizational performance. And yet there is something to be said for getting these issues as close to right as possible—that is, for achieving the second- rather than the third-best (say, by putting the optimal weight on a problematic performance measure in a pay plan, via being fully aware of the costs and benefits of using this measure and this weight). In short, one role for an economist, even in the models summarized above, is to optimize the formal instruments. But will the resulting organization be second-best?

In this section I will argue that economists can do more than optimize the formal instruments and then live with the bleak conclusion of Figure 1 (that organizations live in tough environments, in which performing better than the market would perform is a distinctly backhanded compliment). In particular, I will argue that the key to superior organizational performance is informal: managing relational contracts within and between firms. In making this argument, I am inspired by detailed accounts from observers and practitioners, but I am also conscious that current theory cannot even express (not to mention evaluate) much of what I am trying to say. I hope that in a decade or so, theory will catch up with best practice, enabling someone to write a paper here provisionally titled “What the Folk Theorem Didn’t Tell You.”

In abstract terms, my argument is that conceiving, communicating, and implementing relational contracts appear to be hard enough, but creating, maintaining, and changing relational contracts seem to require real talent and inspiration. There is beginning to be some work in economics along these lines, but delivering big progress will require innovations in the theory, not just applications of existing theory to new questions. Thus, this research area may become a leading example of Kreps’s (199x: yy) observation that “A game-theoretic theory of organizations will do more for game theory than game theory will do for it.”

To set the stage, recall that Coase argued that firms would not need to exist if markets were perfect, and that Williamson (1975) suggested both why markets might perform poorly (because formal contracts are typically incomplete) and why firms might perform better (because firms might use relational contracts). The second piece of Williamson’s argument is clearly correct: firms are riddled with relational contracts;

virtually every collegial and hierarchical relationship in organizations involves important relational contracts. But business relationships are also riddled with relational contracts: supply chains, joint ventures, networks, alliances, and business groups often involve long-run relationships through which the parties reach accommodations when unforeseen or uncontracted-for events occur. Thus, Williamson's (1975) argument seems incomplete: formal contracts may well be imperfect, but as a result we see important relational contracts between firms as well as within.

In this section I discuss the modeling, applications, and future of relational contracts in economics. Section 2.A describes the repeated-game methodology that economists use to analyze relational contracts. In particular, I revisit Kreps's (1990) model of corporate culture outlined in Section 1.B.4 but give much more detail. Section 2.B returns to the Coase-Williamson comparison of markets versus firms, but describes how repeated-game models can allow relational contracts to be taken into account both within and between firms. Finally, Section 2.C suggests that future repeated-game models may shed new light on many structures and processes both within and between organizations, and that such models may even help economists grapple with such unfamiliar topics as management and leadership.

2.A An Introduction to Repeated Games and Relational Contracts¹⁰

Game theory is rampant in economics. Having long ago invaded industrial organization, game-theoretic modeling is now commonplace in international, labor, macro, and public finance, and is gathering steam in development and economic history. Nor is economics alone: accounting, finance, law, marketing, and political science are beginning similar experiences.

Why is this? Broadly speaking, two views are possible: fads and fundamentals. While I believe that fads are partly to blame for the current enthusiasm for game theory, I also believe that fundamentals are an important part of the story. Simply put, many modelers use game theory because it allows them to think like economists when price theory does not apply. Examples abound: small numbers, hidden information, hidden actions, and imperfect contracts can turn markets into games; in other settings, markets are at most peripheral—such as the relationship between a regulator and a firm, a boss and a worker, and so on. Thus, where markets have become games, and where

¹⁰ Those familiar with repeated games and relational contracts may prefer to skip this section. If in doubt as to your mastery of the subject, see whether Figures 2 and 3 below are instantly recognizable.

transactions do not occur in markets, game theory allows economists to study the implications of rationality, self-interest, and equilibrium when price theory would not.

When people interact over time, threats and promises concerning future behavior may influence current behavior. Repeated games capture this fact of life, and hence have been applied more broadly than any other game-theoretic model (by my armchair citation count). In this sub-section I describe first a one-shot interaction between two parties (which works out badly), then an ongoing sequence of such interactions (which work out well because of the parties' concerns for their reputations), and finally how such a repeated-game model captures important aspects of relational contracts in the world at large.

2.A.1 *The One-Shot Interaction*

Suppose that late last night an exciting new project occurred to you. The project would be highly profitable, but is outside your area of expertise, so you would need help in completing it. Furthermore, it would take significant work on your part just to explain the project to someone with the needed expertise. Finally, if you did explain the project to the relevant other, that person could steal your ideas, representing them as substantially his own.

It is not hard to imagine this scenario unfolding in an organization: you work in marketing and the project is a new product, but you need assistance from someone in engineering, who could later take all (or at least too much of) the credit. To decide whether to pursue the project, it would help to know something about the "trustworthiness" of a particular engineer you could approach. But if you have been buried deep inside marketing, you may not have much information about any of the relevant engineers. In this case, you would be forced to rely either on the average sense of human decency among engineers or on your organization's culture—a shared sense of "how we do things around here" (Peters, 19xx: yy). If the culture emphasizes teamwork over individual accomplishments, for example, you may have more confidence in approaching an unfamiliar member of the engineering group.

Kreps (1990) captures these issues in the *Trust Game* in Figure 2. The game begins with a decision node for player 1, who can choose either to Trust or Not Trust player 2. If player 1 chooses Trust then the game reaches a decision node for player 2, who can choose either to Honor or Betray player 1's Trust. If player 1 chooses Not Trust then the game ends. At the end of each branch of the game tree, player 1's payoff appears above player 2's. If player 1 chooses to Not Trust then both players' payoffs are zero. If 1

chooses to Trust 2, however, then both players' payoffs are one if 2 Honors 1's trust, but player 1 receives -1 and player 2 receives two if player 2 Betrays 1's trust.




Figure 2. The Trust Game

We solve the Trust Game by backwards induction—that is, by working backwards through the game tree, one node at a time. If player 2 gets to move (*i.e.*, if player 1 chooses to Trust player 2) then 2 can receive either a payoff of one by Honoring 1's trust or a payoff of two by Betraying 1's trust. Since two exceeds one, player 2 will Betray 1's trust if given the move. Knowing this, player 1's initial choice amounts to either ending the game (and so receiving a payoff of zero) or Trusting player 2 (and so receiving a payoff of -1, after player 2 Betrays 1's Trust). Since zero exceeds -1, player 1 should Not Trust.

2.A.2 The Repeated Game

Instead of a one-shot interaction, suppose that you and a *particular* engineer will play the Trust Game repeatedly, with all previous outcomes observed by both players before the next period's Trust Game is played. The analysis of this repeated game can differ dramatically from the one-shot interaction: the engineer's action today may affect your expectation of her action tomorrow, which may affect your action tomorrow, which affects her payoff tomorrow. Thus, actions not in the engineer's short-run self-interest (as defined by her payoffs today) may be consistent with her overall self-interest (as defined by the present value of her payoffs over time).

I do not mean to imply that this logic is surprising or complicated. To the contrary, I think it is close to ubiquitous. I now describe the simplest possible formalization of this logic. Formally, we will analyze an infinitely repeated game: the game never ends, but both players face an interest rate r per period in discounting their payoffs across periods. (For example, when r is high, a dollar to be received next period is not worth much today— $\$1/(1+r)$, to be exact.) We can interpret this “infinitely” repeated game somewhat more realistically by saying that the game ends at a random date. Under this interpretation, the interest rate r reflects not only the time value of money but also the probability that the relationship will end after the current period. (A dollar to be received next period provided that we are still interacting is not worth much if today’s interaction is likely to be our last.) Under either interpretation, the present value of $\$1$ to be received every period starting tomorrow is $\$1/r$.

Mostly for analytical simplicity (but to some extent for behavioral realism), we will consider the following “trigger” strategies in the infinitely repeated game:

Player 1: In the first period, play Trust. Thereafter, if all moves in all previous periods have been Trust and Honor, play Trust; otherwise, play Not Trust.

Player 2: If given the move this period, play Honor if all moves in all previous periods have been Trust and Honor; otherwise, play Betray.

These strategies are not forgiving, like Tit-for-Tat (see Axelrod, 1984). Rather, under trigger strategies, if cooperation breaks down at any point then it is finished for the rest of the game, replaced by the dictates of short-run self-interest. In most games, reverting to short-run self-interest after a breakdown in cooperation is a middle ground between two plausible alternatives: forgiveness (*i.e.*, an attempt to resuscitate cooperation) and spite (*i.e.*, going against short-run self-interest in order to punish the other player). Both forgiveness and spite deserve analytical attention, but I will focus on the trigger strategies (with their reversion to short-run self-interest after a breakdown of cooperation) as a tractable compromise.¹¹

We will analyze whether these trigger strategies are an equilibrium of the infinitely repeated game. That is, given that player 1 is playing her trigger strategy, is it in player 2’s interest to play his? We will see that the trigger strategies are an equilibrium of

¹¹ In the Trust Game, unlike most, reverting to short-run self-interest is identical to spite: it achieves the harshest possible punishment of player 2.

the infinitely repeated game provided that player 2 is sufficiently patient (*i.e.*, provided that the interest rate r is sufficiently small).

Suppose that player 1 follows his trigger strategy and chooses Trust in the first period. Player 2 then faces a dilemma. As in the one-shot interaction, player 2's one-period payoff is maximized by choosing to Betray. But in the repeated game, if player 1 is playing the trigger strategy then such a betrayal by player 2 leads player 1 to choose No Trust forever after, producing a payoff of zero for player 2 in each subsequent period. Thus, the key question is how player 2 trades off the short-run temptation (a payoff of 2 instead of 1 now) against the long-run cost (a payoff of 0 instead of 1 forever after). The answer depends on the interest rate: if r is sufficiently low then the long-run consideration dominates and player 2 prefers to forego the short-run temptation.

The general point is that cooperation is prone to defection (otherwise we should call cooperation something else—such as a happy alignment of the players' self-interests), but in some circumstances defection can be met with punishment. A potential defector therefore must weigh the present value of continued cooperation against the short-term gain from defection followed by the long-term loss from punishment. If a player's payoffs (per period) are C from cooperation, D from defection, and P from punishment (where $D > C > P$) then this decision amounts to evaluating two time-paths of payoffs: (C, C, C, \dots) versus (D, P, P, P, \dots) , as shown in Figure 3.




Figure 3. Payoffs over time from cooperation or defection today

Because the present value of \$1 received every period starting tomorrow is $\$1/r$, the time-path of payoffs (C, C, C, \dots) yields a higher present value than the time-path (D, P, P, P, \dots) if

$$(*) \quad \left\{ 1 + \frac{1}{r} \right\} C > D + \frac{1}{r} P .$$

Rearranging the inequality (*) yields $r < (C - P)/(D - C)$. In the repeated-games literature, this result is often restated as follows: if the players are sufficiently patient (*i.e.*, if r is sufficiently small) then it is optimal to cooperate, foregoing the short-run temptation ($D - C$ now) for the long-term gain ($C - P$ forever after). For purposes of this chapter, however, it is more useful to recall that the interest rate r reflects not only the time value of money but also the probability that the relationship will end after the current period. If this probability is low then r is low. Thus, the result that cooperation is optimal if r is sufficiently small can be interpreted in terms of something like social structure: provided the time value of money is not too high, cooperation is optimal today if the players' relationship is sufficiently likely to continue in the future.

2.A.3 Repeated-Game Models of Relational Contracts

Both within and between organizations, relational contracts can help circumvent difficulties in formal contracting. For example, a formal contract must be specified *ex ante* in terms that can be verified *ex post* by a third party (such as a court), while a relational contract can be based on outcomes that a court cannot verify *ex post*, and also on outcomes that are prohibitively costly to specify to a court *ex ante*. A relational contract thus allows the parties to utilize their detailed knowledge of their specific situation and to adapt to new information as it becomes available. For the same reasons, however, these relational contracts cannot be enforced by a third party. Instead, relational contracts must be designed to be self-enforcing: each party's reputation must be sufficiently valuable, relative to that party's payoff from renegeing on the relational contract, so that neither party wishes to lose his reputation by renegeing. Put even more abstractly, condition (*) above must hold, where C is the payoff from abiding by the relational contract, D is the payoff from renegeing on the relational contract, and P is the payoff after losing one's reputation for abiding by relational contracts.

The model described here is excessively tidy: cooperation either works perfectly or doesn't work at all, depending on the interest rate. It is natural to ask what happens when the players are not "sufficiently patient." In brief, all is not lost, because it may be possible to achieve partial rather than full cooperation.¹² It is also natural to ask why there

¹² To see how this may work, examine (*). Note that reducing the payoff from cooperation from C to some lower level is no help at all, in and of itself. That is, holding the payoff from defection (D) constant, reducing C makes it harder to satisfy (*). The trick is that reducing C may also reduce D : making due with

are never any fights or misunderstandings in this equilibrium. Green and Porter (1984) develop a model in which the players' actions are not perfectly observable and cooperation breaks down periodically (but for a finite time, after which it begins again). Adding such imperfect observability, and its resulting temporary breakdowns of cooperation, would be a step towards realism in many of the settings considered in the next two sections.

2.B Relational Contracts Within or Between Firms?

Baker, Gibbons, and Murphy (1999) revisit the Coase-Williamson comparison of markets versus firms, taking into account the ubiquity of relational contracts in both domains. In this section I briefly describe our model—not to prove its results, but rather to use the model as an organizing framework for several classic contributions to the theory of the firm.

2.B.1 A One-Shot Supply Transaction

Consider the following model of a one-shot supply transaction involving an upstream party (supplier), a downstream party (user), and an asset (production equipment). Suppose that the upstream party uses the asset to produce a good that can be used in the downstream party's production process. The value of this good to the downstream party is Q , but the good also has an alternative use with value P . Such a supply transaction is shown in Figure 4.

partial cooperation may also limit the players' opportunities for profitable deviations. If D falls more than C then (*) may hold at partial cooperation when it did not at full cooperation.

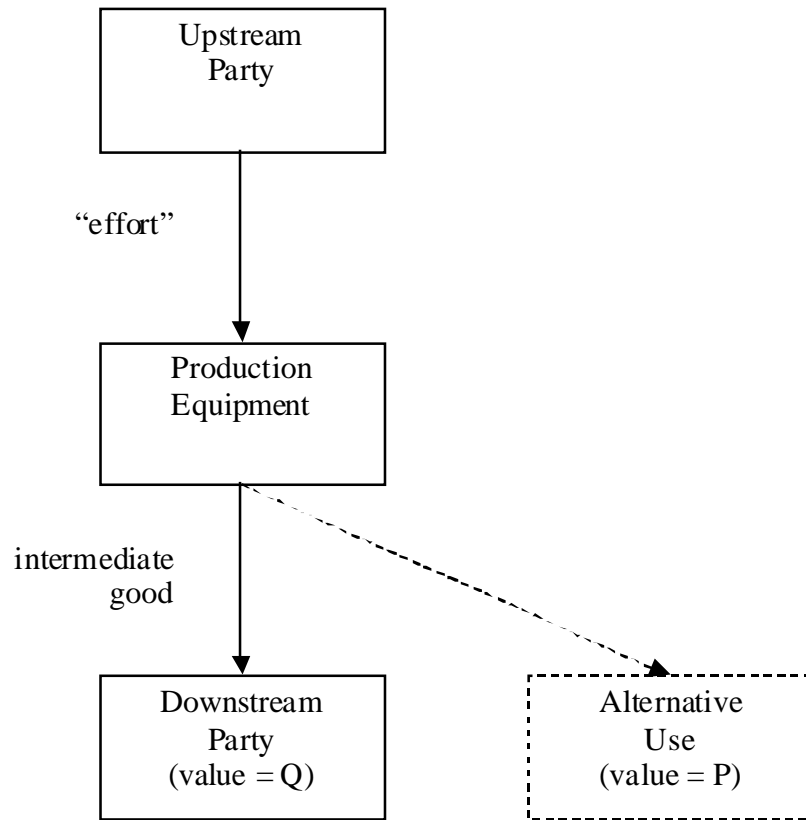


Figure 4. A one-shot supply relationship

If the upstream party owns the asset, call her an *independent contractor* (*i.e.*, someone who works with her own tools); if the downstream party owns the asset, call the upstream party an *employee* of the downstream organization (*i.e.*, someone who works with the boss's tools). Alternatively, one can think of the upstream and downstream parties as firms rather than as individuals, in which case it is more natural to use terms such as *supplier* and *division* rather than independent contractor and employee, respectively. Whether the parties are individuals or firms, if the upstream party owns the asset then call the parties *non-integrated*, but if the downstream party owns the asset then call the parties *integrated*.¹³

To fix ideas, I will cast much of the discussion in terms of a famous business-school case: Crown Cork and Seal Company (Gordon, Reed, and Hamermesh, 1977). The details of the case become important below; for now, it suffices to say that in the 1950s

¹³ Grossman and Hart (1986) originated the idea that asset ownership patterns define who works for whom, and consequently whether a supply chain is integrated or not. As will become clear below, their paper is on a par with Coase (1937) and Williamson (1975, 1985) in its impact on the economic theory of the firm.

and '60s Crown made metal cans for the soft-drink industry. So suppose that Crown owns a can plant located near a Pepsi plant, but there is also a Coke plant two towns away. That is, Crown is the upstream party, Pepsi the downstream party, and Coke the alternative use. In actual fact, Crown was never integrated with Pepsi or Coke, but one can consider the hypothetical case in which Pepsi has purchased the can plant from Crown (in which case the can plant is a "division" of Pepsi).

Suppose that ownership of the asset conveys ownership of the good produced using the asset. For example, if Crown owns the can plant then Crown owns the cans produced there until Pepsi buys them. Furthermore, in bargaining over the sale of the cans, Crown can threaten to sell the cans to Coke (*i.e.*, under non-integration, the upstream party can threaten to consign the good to its alternative use). On the other hand, if Pepsi owned the can plant then Pepsi could prevent the can plant from dealing with outside customers.

Suppose also that the production equipment has been specialized to meet the downstream party's needs. For example, the can plant might have been configured to produce cans to Pepsi's specifications rather than Coke's. Then the good's value to the downstream party will exceed its value in the alternative use; that is, $Q > P$. The surplus that the upstream and downstream parties can jointly achieve by transacting with each other is thus $Q - P$, but each party would like to capture all of this surplus. For example, Crown would like to sell its cans to Pepsi for Q , but Pepsi would like to pay only P .

In Baker, Gibbons, and Murphy (1999) we give the additional details necessary to solve this model of a one-shot supply transaction. In particular, we analyze whether upstream or downstream asset ownership is optimal for a one-shot transaction. But my main interest here is in ongoing supply relationships, so I will skip the details of the one-shot model and move straight to a repeated game.

2.B.2 *An Ongoing Supply Relationship*

In the 1950s and '60s, the metal can industry looked horrible: suppliers were strong (such as U.S. Steel), customers were strong (such as Pepsi, Coke, and Campbell's Soup), and entry into the industry was cheap (a used production line cost only \$150,000 and could be set up in a small space close to an important customer). Industry giants such as American Can and Continental Can were losing money and diversifying out of the industry, but Crown Cork and Seal made money by specializing in customer service in the hard-to-hold segment of the market (*e.g.*, carbonated beverages and aerosol products). That is, Crown began a relationship with a customer by tailoring the specifications of the

cans and the schedule for deliveries to the customer's requirements and detailing these specifications in a formal contract, but (more importantly) Crown also stood ready to modify can specifications and delivery schedules when unusual circumstances arose. Of course, Crown did not make these modifications for free; to the contrary, Crown was able to charge a premium because of its reputation for flexibility and service. In short, in the terminology of this section, Crown had an important relational contract with its customers: Crown would make reasonable modifications under the terms of the existing formal contract; substantial modifications could also be made, but would create the expectation of fair compensation, either on a one-shot basis or by revising the terms of the formal contract for the future.

Crown's customer service illustrates both of Williamson's (1975) ideas. First, formal contracts are almost always incomplete—they often do not specify important future events that might occur, not to mention what adaptations should be made if a particular event does occur. Second, relational contracts may overcome some of the difficulties with formal contracts—relational contracts may allow the parties to utilize their detailed knowledge of their situation to adapt to new contingencies as they arise. Of course, the irony in this illustration is that Crown was not integrated with Pepsi. That is, Crown shows again that relational contracts can be tremendously effective between firms as well as within. A useful model of relational contracts must therefore be applicable both within and between firms.

To analyze relational contracts (within and between firms) in the model in Figure 4, suppose first that the upstream party owns the asset. This case gives rise to the classic "hold-up" problem emphasized by Williamson (1975), because the upstream party can threaten to consign the good to its alternative use unless the downstream party pays a high price. That is, Crown could threaten to sell the cans to Coke. In the model, Pepsi's value for the cans is Q and Coke's is only $P < Q$. Thus, Crown's threat to sell the cans to Coke should not be carried out, because Pepsi is willing to pay more than P for the cans. Instead, after such a threat, suppose that Crown and Pepsi agree on some price between P and Q . The key point is that Crown will receive at least P , and this in turn gives Crown an incentive to take actions that increase P : Crown will pay attention to Coke so as to improve its bargaining position with Pepsi. But actions that increase P may have no (or even negative) effect on Q . Thus, Crown may find it privately optimal to take actions that give it a larger share of a smaller total surplus in its relationship with Pepsi. Such actions are inefficient: both Crown and Pepsi could be made better off if those actions were stopped.

Pepsi's instinctive reaction to this hold-up problem might be the one often prescribed in the transaction-cost literature: buy the can plant, in order to decree that the plant cannot sell cans to Coke. In this sense, vertical integration could indeed prevent one hold-up from occurring, as argued by Williamson (1975) and Klein, Crawford, and Alchian (1978). The insight of Grossman and Hart (1986), however, is that using formal instruments to eliminate one hold-up problem typically creates another. As an example of this conundrum, consider Klein, Crawford, and Alchian's account of the events preceding the acquisition of Fisher Body by General Motors. GM asked Fisher to invest in a new technology to produce closed metal auto bodies rather than the then-standard open wood bodies. Both parties understood that GM could hold-up Fisher after such an investment, such as by offering to pay only marginal rather than average cost. Consequently, the parties signed a contract that gave Fisher certain protections, including a formula specifying the price as a mark-up of Fisher's variable costs. But this contract created ways for Fisher to hold-up GM, such as by threatening to overstaff its plants so as to pad variable cost. Grossman and Hart's abstract model is similar: using asset ownership (another formal instrument, akin to a formal contract) to solve one hold-up problem inevitably creates another.

Ultimately, GM bought Fisher, but at a high price. The price had to be high because Fisher had to be persuaded to give up its strong bargaining position created by the pricing formula in the formal contract. But the reason that it was efficient for GM to buy Fisher does not hinge on this acquisition price, which is merely a transfer between the parties and so has no effect on efficiency considerations. Instead, the reason for GM to buy Fisher (according to Klein, Crawford, and Alchian) was to stop Fisher's inefficient actions, such as overstaffing. Analogously, it might be efficient for Pepsi to buy the can plant from Crown if, under non-integration, Crown has a strong incentive to take inefficient actions that increase the cans' value to Coke (P) but distract Crown from providing service to Pepsi (*i.e.*, reduce Q).

The striking feature of this long-standing and plausible account of the Fisher Body acquisition (see also Klein, 1991) is that it never mentions life in the Fisher division of GM after the acquisition. But without considering the difference between life as a division and life as an independent firm, the analysis cannot ascertain whether the Grossman-Hart conundrum applies. That is, if vertical integration stopped Fisher's hold-up of GM, might it also have created a new way for GM to hold-up Fisher? In keeping with Grossman and Hart, I will argue that integration probably did create such a reverse hold-up. But I will then argue that this conundrum arises because of the reliance on *formal* instruments (such as formal contracts or asset ownership) to eliminate individual

hold-up problems, and that a potential solution to the conundrum is to use *informal* instruments (namely, relational contracts) in tandem with formal instruments to ameliorate all hold-up problems simultaneously. To make these arguments concrete, I return to the Crown-Pepsi example and the model above.

Imagine that Pepsi bought the can plant from Crown. That is, the downstream party owns the asset. The upstream party is then an internal division rather than an external supplier, but the downstream party is still interested in receiving high-quality service. The downstream party could try to create an incentive for the upstream party to supply high-quality service by promising to pay a bonus to the upstream party if the latter produces a sufficiently high value of Q . Unfortunately, like all relational contracts, this promise is vulnerable to renegeing: when the downstream party owns the asset, the downstream party can simply take the intermediate good without paying the upstream party anything.¹⁴

Reneging on a promised bonus is just one example of possible hold-ups within organizations. Richer models could capture renegeing temptations concerning promotions, task allocation, capital allocation, internal auditing transfer payments, and so on. (See Lawler (1971), Bower (1970), Dalton (1959), Eccles (1985), and many others for evidence that such varieties of renegeing are alive and well in many organizations.) The key feature of all of these examples is that one party with authority makes a promise to another party without. In each case, the temptation to renege on such a promise can again be analyzed using Figure 3.

We are now ready to state the key result in this section: the incentive to renege on a relational contact depends on who owns the asset. Suppose the parties would like the upstream party to deliver quality Q^* and the downstream party to pay upstream a fee F^* . Under non-integration, the upstream party is tempted to renege, by taking actions that increase P so as to collect a fee greater than F^* , even if the resulting quality is $Q < Q^*$. Under integration, it is the owner (here, the downstream party) who is tempted to renege, by simply taking the good and not paying the fee F^* . Thus, not only the size of the

¹⁴ In case such renegeing is not immediately plausible, consider the following updated interpretation of the model in Figure 4 – in (say) biotechnology rather than cans. Suppose that the upstream party is an inventor, the downstream party is a manufacturer, and the asset is the inventor's invention. Rather than discuss ownership of a physical asset like a can plant, we now consider ownership of intellectual property – the invention. If the manufacturer will own any inventions that the inventor might produce then the inventor can be thought of as an employee working in the manufacturer's R&D lab. Alternatively, if the inventor will own her inventions then she can sell them either to the manufacturer or to an alternative user. Now imagine that the inventor is an employee in the R&D lab of a large pharmaceutical firm, and suppose the firm has promised to share the profits from inventions 50-50 with the inventor. If the inventor creates a drug worth ten billion dollars, do we expect the firm to keep its promise? How would the situation differ if the inventor had worked in her own independent research firm?

incentive to renege but also the identity of the party tempted to renege depends on who owns the asset.

We therefore have a situation dear to an economist's heart: a tradeoff. Upstream ownership offers the upstream party some recourse should the downstream party renege, and hence decreases the downstream party's temptation to renege, but upstream ownership also encourages the upstream party to consider the interests of other parties, and hence may create a temptation for the upstream party to renege. In some settings, the first of these considerations is more important, so integration is optimal; in others, the second dominates, so non-integration is preferred. In all settings, however, the guiding principle is to induce efficient actions (and discourage inefficient actions) by implementing the best possible relational contract. The integration decision is merely an instrument in this quest.

2.B.3 Organizing the Classics (and Beyond)

Thus far in Section 2.B I hope to have motivated three assertions: (1) there are many relational contracts that have important influences on behavior, both within and between firms; (2) repeated-game models can capture some of the importance features and tensions of these relational contracts; and (3) the economic incentive to renege on a given relational contract depends on whether the parties are integrated or not. Of course, not all transactions occur in ongoing relationships. Some occur in one-shot relationships (perhaps because the parties cannot meet again, or perhaps because the transaction can be conducted efficiently without the benefit of a relational contract). Thus, one can imagine two polar cases: a "spot" social structure, in which the parties will never see each other again after their current interaction, and a "relational" social structure, in which the parties are sure to interact for the foreseeable future.

These two polar social structures (spot and relational) can be combined with the two possibilities for asset ownership described above (integrated or not), producing the four governance regimes summarized in Figure 5. Consistent with common usage, one can call the integrated case "employment" (where the upstream party is an employee, working with an asset owned by the firm) and the non-integrated case "outsourcing" (where the upstream party is an independent contractor, working with an asset she owns). The four governance regimes are then spot and relational outsourcing and spot and relational employment.

Figure 5

Combinations of social structure and asset ownership that define four governance regimes: Spot Outsourcing, Spot Employment, Relational Outsourcing, and Relational Employment

| | | Asset Ownership | |
|------------------|------------|------------------------|-----------------------|
| | | Non-Integrated | Integrated |
| Social Structure | Spot | Spot Outsourcing | Spot Employment |
| | Relational | Relational Outsourcing | Relational Employment |

Many of the classic contributions to organizational economics can be described using Figure 5. For example, static analyses of integration in the absence of relational contracting (*e.g.*, Grossman and Hart, 1986) are analogous to comparing spot outsourcing to spot employment (the top row in the figure). Similarly, repeated-game analyses of relationships within firms (*e.g.*, Kreps, 1990) are analogous to comparing spot employment to relational employment (the right column) and repeated-game analyses of relationships between firms (*e.g.*, Klein and Leffler, 1981) are analogous to comparing spot outsourcing to relational outsourcing (the left column). Finally, Williamson’s (1975) argument that the comparative advantage of firms over markets lies in the firm’s ability to enforce relational contracts is analogous to comparing spot outsourcing to relational employment (the main diagonal).¹⁵

Baker, Murphy, and I analyze the efficiency of the four governance regimes in Figure 5. In fact, we consider all possible values of the interest rate r , and hence consider all the social structures between the spot and relational polar cases just defined. For any given value of r , we compute whether integration or non-integration is more efficient (*i.e.*, generates greater total surplus for the two parties). In other words, for any given probability that the parties’ relationship will continue in the future, we conduct a formal version of the Coase-Williamson horse race between markets and firms, taking relational contracts into account in both domains.¹⁶

¹⁵ Williamson (1985) emphasized relational contracts between firms, but construed them as lying on a continuum between markets and hierarchies (p. 83). Figure 5 suggests that the set of alternative governance regimes is two-dimensional, so it is not possible to locate all governance structures on a line between markets and hierarchies. Compared to Williamson’s first book, the second showed much greater concern for limits on the effectiveness of firms and the importance of relational contracts between firms. But there was little analysis of how relational contracts between firms might differ from those within, and discussion of alternative organizational forms was limited to brief mention of “hybrid transactions (franchising, joint ventures, and other forms of nonstandard contracting)” (p. 83).

¹⁶ Williamson (1996, Chapter 4) devotes even greater attention than (1985, Chapter 3) to relational contracts between firms, recognizing the prevalence and longevity of these relationships and their

I now hope to have motivated a fourth assertion, in addition to the three above: repeated-game models can be used to run the Coase-Williamson horse race between markets and firms, taking relational contracts into account in both domains. The result is a new perspective on vertical integration, in which the “make or buy” decision is seen as an instrument in the service of the parties’ relationship.

2.C Future Research

Baker, Murphy, and I offer some simple and preliminary suggestions for how our basic repeated-game model could be enriched in future work to analyze specific relational contracts within organizations (such as transfer pricing and capital allocation) as well as the roles of relational contracts in specific multi-organizational forms (such as networks and joint ventures). But we do not offer any suggestions for modeling the creation, maintenance, or change of a repeated-game equilibrium. I believe that these latter issues are central to management and leadership, and that coming to grips with them in a formal model will (in something like Kreps’s words) do as much for game theory as for the theory of organizations.

2.C.1 *Formal and Informal Organizational Structures and Processes*

Organizational sociologists have long emphasized the distinction between formal and informal aspects of organizational structure. Formal aspects include official job descriptions and reporting relationships, as well as formal contracts. Informal aspects include norms and mutual understandings, as well as networks of non-reporting relationships among individuals. Roughly speaking, the formal structure is the organization chart, whereas the informal structure is the way things really work.

In the abstract model in Section 2.B, asset ownership is a formal aspect of organizational structure, relational contracts are informal. The model shows that asset ownership and relational contracts interact, rather than just co-exist. But this conclusion merely echoes Blau and Scott’s (1962: 6) claim from Section 1.B.4 that “... the formally instituted and the informal emerging patterns are inextricably intertwined.” Having developed a formal model, however, we can push towards further conclusions. For example, we analyzed how to choose asset ownership to facilitate relational contracting,

associated “hybrid” organizational forms. The contributions of the Baker-Gibbons-Murphy model are to ask and answer how and why relational contracts between firms might differ from those within, and consequently when one might outperform the other.

thereby optimizing overall organizational performance. And in richer models, one could pursue the same idea more generally: choose formal structure to facilitate informal structure.

I interpret a relational contract between non-integrated parties in the repeated-game model in Section 2.B as a hand-in-glove supply relationship. Thus, the existing model does speak to one organizational form beyond markets and hierarchies. But there are now many organizational forms discussed in the business and organizational literatures, including joint ventures, strategic alliances, networks, and business groups. It seems likely that informal aspects, especially relational contracts, are important to the success of these organizational forms. It also seems likely that the formal and informal aspects again interact (rather than just co-exist), so one could choose the former to facilitate the latter. Although the model above has only two stages of production with one party at each stage, richer models could add both parties and stages. For example, one could begin to model a joint venture as two parties at one stage who create an asset at the other stage that they control by both formal and informal means. Similarly, one could begin to model a business group as several parties at several stages of production, with both cross-ownership and relational contracts linking the parties, possibly through a central party. Formal structures such as fifty-fifty ownership in joint ventures or minority stock holdings in business groups may be better understood using models that study the interplay between these formal structures and the informal relational contracts between the parties.

Of course, the interplay between formal and informal organizational structures is not limited to relationships between firms; there are several applications of this idea within firms. Subjective judgments, and consequently parties' reputations, play crucial roles in organizational processes such as transfer pricing, capital allocation, internal auditing, and compensation. That is, many formal organizational processes have important informal aspects; they are not simple exercises in formal contracting, as they often appear in the economics literature. It seems likely that these informal aspects can be hampered or facilitated by formal organizational structures, such as whether the firm is organized by function versus division, whether its jobs are defined for individuals or teams, and so on. Thus, one can again envision a class of models in which formal organizational structures interact with informal relational contracts. Of course, informal relational contracts often hurt rather than improve organizational performance, as famously described by Roethlisberger and Dickson (1939), Roy (1952), and many others. It would be very interesting if certain formal structures were better able than others to facilitate useful rather than destructive relational contracts.

2.C.2 Management and the Nature of the Firm

The literature on the theory of the firm makes little reference to managers. By emphasizing the importance of relational contracts, our model highlights a role for managers: conceiving, communicating, and implementing relational contracts, both in firms and in supplier relationships. In our model, conceiving of the relational contract means determining what the observable but non-contractible Q 's and P 's should be, and determining the optimal b 's and β 's); this conception must then be communicated to employees; finally, to implement the relational contract, the manager must assess the outcomes (that is, observe the Q 's and P 's) and decide whether to honor the relational contract (that is, evaluate the reneging constraints). In our model these tasks are trivial; in reality each requires judgment and knowledge of the specifics of complex situations.

We are not arguing that management occurs only in firms. To the contrary, we have emphasized that management is crucial in ongoing supply relationships (the lower-left box in Figure 5: relational outsourcing). Nor are we arguing that firms cannot occur without management. We can imagine firms that fit the Grossman-Hart model (the upper-right box: spot employment), in which all non-contractible rights of control are held by a single individual. But if these rights become too numerous and complex to be held by a single individual, the firm must have some mechanism to coordinate their exercise, and because these rights are not contractible, their exercise cannot be controlled by formal contracts. Rather, the exercise of non-contractible rights must be coordinated by relational contracts. In short, the firm must have managers. We therefore reach a conclusion that is either obvious or powerful: understanding the role of managers, who conceive, communicate, and implement the relational contracts that underpin informal organizational processes, is essential to understanding firms.

2.C.3 Leadership

Rogers and Beer (1995a, b) describe how Hewlett-Packard (HP) sought to change its clear and strong promise to its employees that there would be no layoffs from HP. For the sake of brief argument, imagine that there was a clear and compelling business case that HP must abandon its no-layoffs policy. One approach would be simply to announce layoffs, thus reneging on the promise. Another would be to engage in detailed communications explaining the necessity of layoffs, and then implement layoffs unilaterally, as in the first approach. A third approach, taken by HP, would be to engage in detailed and persistent discussions about the causes and consequences of the layoff

decision, and then do everything possible to encourage consensus to emerge concerning the layoff issue.

In the language of this section, HP had one relational contract with its employees (no layoffs) but sought to establish another. It would be easy to renege on the original relational contract, but probably then hard to establish the next. Thus, HP faced a problem in changing from one equilibrium to another, not simply in departing from the original equilibrium. I know nothing in the theory of repeated games that helps with this question. In fact, I know nothing that helps with the parallel question of how one builds an equilibrium in the first place. Thus, the managerial tasks of conceiving, communicating, and implementing a relational contract may be hard, but creating, maintaining, and changing an equilibrium may require real talent and inspiration. Again, in a decade or so, I hope someone writes a paper on these issues, explaining “What the Folk Theorem Didn’t Tell You.”

3. Outline, Limitations, and Value of the Book

In this chapter I have tried to convey the flavor of the book that follows—briefly, organizations are a mess but not a mystery, and managing relational contracts may offer a way out. This section first outlines the remainder of the book, showing that many of the book’s six parts expand on the models summarized in Sections 1.B and 2.B above, and then begins what I hope will be an ongoing discussion concerning the limitations of an economic approach to organizations.

3.A Outline of the Book

This book tries to cover a lot of ground. For model building, I draw on agency theory, contract theory, and game theory. For some of the evidence, I draw on labor economics, industrial organization, and economic development. And to broaden my perspectives on both theory and evidence, I try to tap human resource management, organizational psychology, organizational sociology, and parts of political science. Unfortunately, I am not capable of providing much depth across such breadth. But I feel strongly that the emerging field of organizational economics has and needs theoretical and empirical inputs from all these sources, so I have chosen to include at least a bit from each, and consequently cannot cover any exhaustively. Apologies are hereby tendered to

those whose work has been short-changed or omitted. I would be happy to receive suggestions for a second edition.

After this introductory chapter, the book is organized into six parts, followed by a concluding chapter. Part One concerns incentives in organizations. The first chapter considers pay for performance, in the spirit of the “get what you pay for” models in Section 1.B.1; the second discusses subjective performance assessments, in the spirit of Section I.B.4. A third chapter on static models analyzes the tradeoff between insurance and incentives and incentives in groups, and a fourth on dynamic models covers the ratchet effect and career concerns.

Part Two concerns the boundaries of organizations. The first chapter summarizes the literature on the classic “make or buy” problem, in both its original Coase-Williamson form as well as recent incarnations, starting with Grossman-Hart. The second and third chapters consider contracts between firms—first formal contracts, then relational contracts. The final chapter discusses structures between organizations, such as networks, alliances, and joint ventures. Much of the material in these four chapters can be organized using Figure 5 in Section 2.B.3.

Part Three concerns careers in organizations. The first three chapters analyze different aspects of job assignments and promotions: job assignments as a response to differences in workers’ skills; tournaments, in which the prospect of a new job assignment (and a higher wage) creates incentives; and skill development, including incentives for investment in human capital from promotion rules like those discussed in Section 1.B.2. Other chapters analyze integrative models (that attempt to explain broad patterns of evidence on careers in organizations), new work practices (such as self-managed teams), and networks within organizations.

Part Four concerns life in organizations. The first chapter considers communication and decision-making in organizations, ranging from Marschak and Radner’s team theory to Cohen, March, and Olsen’s garbage can, as discussed in Section 1.A. The second chapter covers power and politics, similar to the discussion of influence activities in Section 1.B.3 and March and Pfeffer’s work discussed in Section 1.B.6. The final chapter analyzes culture and leadership, in the spirit of the repeated-game model discussed in Sections 1.B.4 and 2.C.3.

Part Five concerns organizational structures and processes. The first chapter discusses the transition from functional to multi-divisional organizations, starting with Chandler’s (1962) classic work and including Williamson’s (1985) analysis of this issue using the principles of transaction-cost economics. The second chapter considers first formal models of hierarchy and then qualitative accounts of alternative organizational

structures, including cooperatives, partnerships, matrices, and networks. The final chapter covers several organizational processes, including capital budgeting and transfer pricing.

Part Six concerns institutions other than firms. Since the material in this part may be unfamiliar to many readers, I will describe here it in a bit more detail. The first chapter discusses government agencies, which can be seen as coming into existence only in problematic environments (akin to those that spawn firms in Figure 1 above). As a result, government agencies are also a mess but not a mystery. In one of the classic books on government agencies, Wilson (1989, Chapter 9) discusses the agency's compliance with its prescribed methods and goals, and uses terms much like those used to discuss pay for performance and subjective performance assessment in Sections 1.B.1 and 1.B.4. More generally, Moe (1990, 1997) has long been calling for and working towards a productive relationship between organizational economics and the study of government agencies.

The second chapter in Part Six summarizes recent work in economic development and political economy on "order without law" (*e.g.*, among the oligarchs in 13th-century Genoa; see Greif, 1994). This work involves (formal and informal) repeated-game models similar to those of relational contracts between firms discussed in Section 2.B.2. The literature on order without law shows both the flexibility and the empirical purchase of this repeated-game approach. The final chapter reverses the emphasis by considering "law without order:" formal and informal repeated-game models of states as rulers who can extract resources from their citizens (*e.g.*, by raising taxes) but may find it optimal to restrain themselves (see, for example, North and Weingast, 1989).

After these six parts, the concluding chapter summarizes overarching themes and suggests directions for future research.

3.B Limits of an Economic Approach to Organizations

While I have complained that this book is in some respects too broad for me to handle, I am also keenly aware that the economic approach I emphasize is also quite narrow, or even distorted. Examples abound. Where I assume self-interested preferences, reality may be social utility (such as where I care about your payoff; Kelley and Thibaut, 1978; Camerer, 1997). Where I assume exogenous preferences, reality may be contingent preferences (such as where how I feel about your payoff depends on how I think you feel about mine; Rabin, 1993). Turning from preferences to cognition, where I assume full rationality, reality may involve biases and heuristics (Tversky and Kahneman, 1974; Babcock and Loewenstein, 1997) and attribution errors (Ross, 1977; Weber et. al. 1999). And moving from psychology to sociology, where I examine actions in a fixed context (of

networks, institutions, and shared histories), reality may involve not only shifts in context but feedback from actions to contexts (Giddens, 1984; Granovetter, 1999).

By citing some recent work in this connection, some of which is even by economists, I hope to suggest that these issues are slowly moving onto some economists' radar screens and may be incorporated into models in organizational economics in the finite future. Here, as elsewhere, I am all ears.

3.C Virtues of an Economic Approach to Organizations

As of this writing, theory has run far ahead of evidence in the emerging field of organizational economics. Consequently, a very high fraction of this book is devoted to theoretical research, but I hope no one will read me as arguing that formal modeling should be the only style of research in organizational economics. To the contrary, I think the most successful literatures are those that blend detailed description, informal theory, and formal modeling. Against this standard, I think organizational economics has too little description and informal theory, but organizational sociology too little formal modeling.¹⁷ I hope that the efforts in this book to both summarize the theory and identify the existing evidence may help prompt more empirical work.

Fortunately, at least some of the theoretical models in this book were not constructed in an empirical vacuum. Instead, these theories were developed in order to explain an existing (or even alleged) fact. Of course, this research strategy brings its own difficulties. To put one version of this point quite starkly, is there any value in using economic models to formalize ideas that sociologists articulated decades before? I see four answers, versions of which apply to formal models of any kind, not just the rational-choice models familiar to economists.

First, formal rational-choice models check the internal consistency of informal rational-choice arguments. Of course, far from all arguments in organizational sociology are informal rational-choice arguments, but some are. There is value in formalizing such informal rational-choice arguments, because seemingly airtight informal arguments are

¹⁷ Recent work on the organizations and institutions of government may be getting the blend about right. For example, on limited government, see North and Weingast (1989) for description and informal theory and Weingast (1995, 1997) for formal models inspired by detailed descriptions, and on government agencies, see Wilson (1989) for description and informal theory and Moe (1997) for a survey of informal theory and formal models.

sometimes wrong and are sometimes right only if additional assumptions are imposed.¹⁸ Often, the additional assumptions exposed by the formal modeling can be viewed as boundary conditions for an informal argument that is correct except for its omission of such conditions.

Second, formal models can help to specify and interpret empirical tests. This point is related to the first, but here the potentially troublesome step is the informal translation of a theoretical prediction into an empirical test, regardless of whether the theoretical prediction is derived from a formal or an informal theory. For example, consider Granovetter's (1973) argument that weak ties relay job offers more frequently than do strong ties, or Lin's (1982) argument that weak-tie offers are drawn from a better distribution of wages. Bridges and Villemez (1986) found that, after controlling for worker characteristics, tie strength has no effect on wages, seemingly rejecting Granovetter's and Lin's theories. But Montgomery (1992) developed a formal model of network composition showing that Granovetter's and Lin's theories do not imply that observed wages are higher for those who found their jobs through weak ties. That is, Montgomery's formal model showed that a seemingly intuitive empirical test was not an implication of Granovetter's and Lin's theories.¹⁹

¹⁸ A few readers of this essay have asked for examples of incorrect informal arguments—perhaps because such examples would improve their understanding of my argument, but perhaps also because such examples make good gossip. The infrequency of formal modeling in organizational sociology makes examples hard to generate in that field, so I will stick to organizational economics. I will also choose authors whose reputations are sufficiently strong that being mentioned in this context will do no harm. One example involves career concerns—that is, a worker's concerns about the effects of his or her current performance on future compensation and employment opportunities. Career concerns were first discussed by Fama (1980), who gave an informal argument that managerial incentive contracts are not necessary, because managers are disciplined through the managerial labor market: superior performances will generate high wage offers; poor performances, low offers. Holmstrom (1982b) then developed a formal model that captured Fama's ideas about labor-market discipline but showed that such discipline is nothing like a perfect substitute for contracts: in the absence of contracts, managers typically work too hard in early years (while the market is still assessing the manager's ability) and not hard enough in later years. I give another example in the next footnote, as part of the discussion of how formal models can help to specify and interpret empirical tests.

¹⁹ As a second example, consider Spence's (1973) signaling model of education, in which workers choose education levels in an attempt to signal their productive abilities to potential employers. In this model, firms rationally offer higher wages to those who acquire more education, even if education has no effect on a worker's productivity. Taken literally, Spence's model deals only with the wage paid at the beginning of a worker's first job after school, but Layard and Psacharopoulos (1974) suggested that the model could be tested by studying the effect of education on a worker's wages at various levels of labor-market experience, as follows. If firms observe some indication of a worker's performance while he or she is in the labor market, firms should use these indications to update their beliefs about the worker's ability, and hence to update their wage offers to the worker. It would therefore seem that the effect of education on

Third, some formal rational-choice models may provide new explanations for facts previously given non-rational-choice interpretations. For example, Banerjee's (1992) rational-choice model of herd behavior (summarized in Section 1.B.5) has close parallels to institutionalist accounts of mimetic isomorphism (DiMaggio and Powell, 1983). In an ideal world, the presence of competing explanations for a given set of facts would inspire empirical testing of the theories' auxiliary implications.

Finally, developing formal models in rich contexts sometimes allows the analysis to push through to conclusions when informal analysis would have ground to a halt. For instance, even if two informal arguments are both airtight and simple, it may be difficult to think informally about how the arguments interact, yet their interaction may produce new insights or predictions. The eminent mathematical sociologist James Coleman (1964: vii) made this point much more eloquently: "If conceptual elaboration is to progress beyond the proverbs of the ancients, special tools are necessary. The most remarkable of these is mathematics. ... The mind falters when faced with a complex system or a long chain of deductions. The crutch that mathematics provides to everyday reasoning becomes essential as sociology moves toward the analysis of complex systems and predictions based on extended chains of deductions."

For all these reasons, I find the theories summarized in this book useful, even when they are formalizing earlier informal arguments or offering explanations for existing facts. If organizational economics is to become a full-fledged field, it will of course be crucial to develop more evidence and to test the auxiliary implications of existing and future models. But neither seems likely to happen without first summarizing what we have learned so far.

wages should decrease with labor-market experience. But Farber and Gibbons (1996) developed a formal model of such learning in labor markets and showed that, while education does come to play a smaller role in firms' beliefs, the estimated effect of education on wages in a linear regression should be constant at all levels of labor-market experience, because of omitted-variable bias. Thus, once again, a seemingly intuitive but informally derived empirical test was not in fact an implication of the theory in question. Furthermore, in this case, the formal model delivered an empirical test that is an implication of the theory and that is not rejected by the data.

References

- Akerlof, George. 1970. "The market for lemons: Quality uncertainty and the market mechanism." *Quarterly Journal of Economics* 89: 488-500.
- Arrow, Kenneth and Frank Hahn. 1972. *General Competitive Analysis*. Amsterdam: North-Holland Publishing Company.
- Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Babcock, L., and Loewenstein, G. 1997. "Explaining Bargaining Impasse: The Role of Self-Serving Biases." *Journal of Economic Perspectives* 11: 109-26.
- Baker, George. 1992. "Incentive Contracts and Performance Measurement." *Journal of Political Economy* 100:598-614.
- _____, Robert Gibbons, and Kevin J. Murphy. 1994. "Subjective Performance Measures in Optimal Incentive Contracts." *Quarterly Journal of Economics* 109:1125-56.
- _____, _____, and _____. 1999. "Relational contracts and the theory of the firm." Unpublished manuscript, MIT's Sloan School.
- Banerjee, Abhijit. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107:797-817.
- Barnard, Chester. 1938. *The Functions of the Executive*. Cambridge, MA: Harvard University Press.
- Becker, Gary. 1964. *Human Capital*. New York: NBER.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100: 992-1026.
- _____, _____, _____. 1998. "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades." *Journal of Economic Perspectives* 12: 151-70.
- Blau, Peter. 1955. *The Dynamics of Bureaucracy*. Chicago: University of Chicago Press.
- _____, and W. Richard Scott. 1962. *Formal Organizations: A Comparative Approach*. San Francisco: Chandler.
- Blumenstein, Rebecca and Gabriella Stern. 1996. "UAW, Auto Makers Find Some Things Better Left Unsaid." *Wall Street Journal*, November 25, A1.
- Bower, Joseph. 1970. *Managing the Resource Allocation Process*. Boston: Harvard Business School Press.
- Bridges, William, and Wayne Villemez. 1986. "Informal hiring and income in the labor market." *American Sociological Review* 51: 574-582.
- Bull, Clive. 1987. "The Existence of Self-Enforcing Implicit Contracts," *Quarterly Journal of Economics* 102:147-59.

- Camerer, Colin. 1997. "Progress in Behavioral Game Theory." *Journal of Economic Perspectives* 11: 167-88.
- Chandler, Alfred. 1962. *Strategy and Structure*. Cambridge, MA: MIT Press.
- Coase, Ronald. 1937. "The nature of the firm." *Economica* 4: 386-405.
- Cohen, Michael, James March, and Johan Olsen. 1972. "A Garbage Can Model of Organizational Choice." *Administrative Science Quarterly* 17:1-25.
- Coleman, James. 1964. *Introduction to Mathematical Sociology*. London: Collier-Macmillan Limited.
- Crozier, Michel. 1964. *The Bureaucratic Phenomenon*. Chicago: University of Chicago Press.
- Dalton, Melville. 1959. *Men Who Manage*. New York: Wiley.
- DiMaggio, Paul, and Walter Powell. 1983. "The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields." *American Sociological Review* 48: 147-60.
- _____ and _____. 1991. "Introduction." In W. Powell and P. DiMaggio (eds.), *The New Institutionalism in Organizational Analysis*. Chicago: University of Chicago Press.
- Dore, Ronald. 1983. "Goodwill and the spirit of market capitalism." *British Journal of Sociology* 34: 459-482.
- Eccles, Robert. 1981. "The quasifirm in the construction industry." *Journal of Economic Behavior and Organization* 2: 335-357.
- _____. 1985. *The Transfer Pricing Problem: A Theory for Practice*. Lexington, MA: D.C. Heath.
- Fama, Eugene. 1980. "Agency problems and the theory of the firm." *Journal of Political Economy* 88: 288-307.
- Farber, Henry and Robert Gibbons. 1996. "Learning and wage dynamics." *Quarterly Journal of Economics* 111: 1007-47.
- Feldman, Martha and James March. 1981. "Information in Organizations as Signal and Symbol." *Administrative Science Quarterly* 26:171-86.
- Fudenberg, Drew, and Eric Maskin. 1986. "The folk theorem in repeated games with discounting or with incomplete information." *Econometrica* 54: 532-554.
- Gibbons, Robert. 1992. *Game Theory for Applied Economists*. Princeton, NJ: Princeton University Press.
- _____. 1997. "An Introduction to Game Theory for Applied Economists." *Journal of Economic Perspectives* 11: 127-49.
- _____. 1998a. "Game theory and garbage cans: An introduction to the economics of internal organization." In J. Halpern and R. Stern (eds.), *Nonrational Elements of Organizational Decision Making*. Ithaca, NY: ILR Press.
- _____. 1998b. "Incentives in organizations." *Journal of Economic Perspectives* 12: 115-132.
- _____. 1999. "Taking Coase Seriously." *Administrative Science Quarterly* 44: 145-57.

- _____. 2000. "Firms (and Other Relationships)." Forthcoming in P. DiMaggio (ed.), *The Twenty-First Century Firm: Changing Economic Organization in International Perspective*. Princeton, NJ: Princeton University Press.
- Giddens, Anthony. 1984. *The Constitution of Society*. Berkeley, CA: The University of California Press.
- Gordon, Karen, John Reed, and Richard Hamermesh. 1977. "Crown Cork and Seal Company, Inc." Harvard Business School Case #378-024.
- Gouldner, Alvin. 1954. *Patterns of Industrial Bureaucracy*. New York: Free Press.
- Granovetter, Mark. 1973. "The strength of weak ties." *American Journal of Sociology* 78: 1360-80.
- _____. 1985. "Economic Action and Social Structure: The Problem of Embeddedness." *American Journal of Sociology* 91:481-510.
- _____. 1999. "Coase Encounters and Formal Models: Taking Gibbons Seriously." *Administrative Science Quarterly* 44: 158-62.
- Green, Edward, and Robert Porter. 1984. "Noncooperative Collusion Under Imperfect Price Information." *Econometrica* 52: 87-100.
- Greif, Avner. 1994. "On the Political Foundations of the Late Medieval Commercial Revolution: Genoa During the Twelfth and Thirteenth Centuries." *Journal of Economic History* 54: 271-87.
- Grossman, Sanford, and Oliver Hart. 1986. "The costs and benefits of ownership: A theory of vertical and lateral integration." *Journal of Political Economy* 94: 691-719.
- Hart, Oliver. 1995. *Firms, Contracts, and Financial Structure*. Oxford: Oxford University Press.
- _____ and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy* 98: 1119-58.
- Heckman, James. 1976. "The common structure of statistical models of truncation, sample selection, and limited dependent variables and a simple estimator for such models." *Annals of Economic and Social Measurement* 5: 475-92.
- Holmstrom, Bengt. 1979. "Moral hazard and observability." *Bell Journal of Economics* 10: 74-91.
- _____. 1982a. "Moral hazard in teams." *Bell Journal of Economics* 13: 324-340.
- _____. 1982b. "Managerial Incentive Problems—A Dynamic Perspective." In *Essays in Economics and Management in Honor of Lars Wahlbeck*. Helsinki: Swedish School of Economics. Republished in *Review of Economic Studies* 66 (1999): 169-82.
- _____ and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, and Organization* 7:24-52.
- Kahn, Charles and Gur Huberman. 1988. "Two-sided Uncertainty and 'Up-or-Out' Contracts." *Journal of Labor Economics*. 6:423-44.
- Kelley, H. and J. Thibaut. 1978. *Interpersonal Relations: A Theory of Mixed Interdependence*. New York: John Wiley & Sons.
- Kerr, Steven. 1975. "On the Folly of Rewarding A, While Hoping for B." *Academy of Management Journal* 18:769-83.

- Klein, Benjamin. 1991. "Vertical Integration as Organizational Ownership: The Fisher Body-General Motors Relationship Revisited." In O. Williamson and S. Winter (eds.), *The Nature of the Firm: Origins, Evolution, and Development*. Oxford University Press.
- _____, Robert Crawford, and Armen Alchian. 1978. "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process." *Journal of Law and Economics* 21: 297-326.
- _____ and Keith Leffler. 1981. "The Role of Market Forces in Assuring Contractual Performance." *Journal of Political Economy*. 89:615-641.
- Kreps, David. 1990. "Corporate culture and economic theory." In J. Alt and K. Shepsle (eds.), *Perspectives on Positive Political Economy*. New York: Cambridge University Press.
- Kreps, David. 199x.
- _____, Paul Milgrom, John Roberts, and Robert Wilson. 1982. "Rational cooperation in the finitely repeated prisoners' dilemma." *Journal of Economic Theory* 27: 245-252.
- Lawler, Edward. 1971. *Pay and Organizational Effectiveness: A Psychological View*. New York: McGraw-Hill.
- Layard, Richard and George Psacharopoulos. 1974. "The screening hypothesis and the returns to education." *Journal of Political Economy* 82: 985-98.
- Lin, Nan. 1982. "Social resources and instrumental action." In P. Marsden and N. Lin (eds.), *Social Structure and Network Analysis*. Sage Publications: Beverly Hills, CA.
- Macaulay, Stewart. 1963. "Non-contractual relations in business: A preliminary study." *American Sociological Review* 28: 55-67.
- March, James. 1962. "The business firm as a political coalition." *Journal of Politics* 24: 662-678.
- March, James. 1994. *A Primer on Decision Making: How Decisions Happen*. New York: Free Press.
- Marschak, Jacob, and Roy Radner. 1972. *Economic Theory of Teams*. New Haven, CT: Yale University Press.
- Merton, Robert. 1940. "Bureaucratic structure and personality." *Social Forces* 18: 560-568.
- Meyer, John, and Brian Rowan. 1977. "Institutionalized organizations: Formal structure as myth and ceremony." *American Journal of Sociology* 83: 340-363.
- Milgrom, Paul and John Roberts. 1988. "An Economic Approach to Influence Activities in Organizations." *American Journal of Sociology* 94:S154-S179.
- Moe, Terry. 1990.
- _____. 1997. "The Positive Theory of Public Bureaucracy." In D. Mueller (ed.), *Perspectives on Public Choice: A Handbook*. New York: Cambridge University Press.
- Montgomery, James. 1992. "Job search and network composition: Implications of the strength-of-weak-ties hypothesis." *American Sociological Review* 57: 586-96.
- North, Douglas and Barry Weingast. 1989. "Constitutions and Commitment: The Evolution of Institutions Governing Public Choice in Seventeenth-Century England." *The Journal of Economic History* 49:803-32.

Peters, Tom.

Pfeffer, Jeffrey. 1981. *Power in Organizations*. New York: HarperBusiness.

Podolny, Joel, and Karen Page. 1998. "Network Forms of Organization." *Annual Review of Sociology*, 24: 57-76.

Post, Richard J. and Kenneth E. Goodpaster. 1981. "H.J. Heinz Company: The Administration of Policy." Harvard Business School Case #382-034.

Powell, Walter. 1990. "Neither Market nor Hierarchy: Network Forms of Organization." *Research in Organizational Behavior*. 12:295-336.

_____. 1991. "Expanding the Scope of Institutional Analysis." Chapter 8 in W. Powell and P. DiMaggio (eds.), *The New Institutionalism in Organizational Analysis*. Chicago: University of Chicago Press.

Prendergast, Canice. 1993a. "The Role of Promotion in Inducing Specific Human Capital Acquisition." *Quarterly Journal of Economics* 108:523-34.

_____. 1993b. "A Theory of 'Yes Men.'" *American Economic Review* 83:757-70.

Rabin, M. 1993. "Incorporating fairness into game theory and economics." *American Economic Review* 83: 1408-18.

_____. 1998. "Psychology and Economics." *Journal of Economic Literature* 36: 11-46.

Rajan, Raguram, and Luigi Zingales. 1996. "The tyranny of the inefficient: An enquiry into the adverse consequences of power struggles." Unpublished manuscript, Graduate School of Business, University of Chicago.

Roethlisberger, Fritz, and William Dickson. 1939. *Management and the Worker*. Cambridge: Harvard University Press.

Rogers, Gregory and Michael Beer. 1995a. "Human Resources at Hewlett-Packard (A)." Harvard Business School Case #9-495-051.

Rogers, Gregory and Michael Beer. 1995b. "Human Resources at Hewlett-Packard (B)." Harvard Business School Case #9-495-052.

Ross, L. 1977. "The intuitive psychologist and his shortcomings." In L. Berkowitz (Ed.), *Advances in experimental social psychology*, (Volume 10). San Diego, CA: Academic Press.

Rotemberg, Julio. 1994. "Power in profit-maximizing organizations." *Journal of Economics and Management Strategy* 2: 165-198.

Rothschild, Michael, and Joseph Stiglitz. 1976. "Equilibrium in competitive insurance markets: An essay on the economics of imperfect information." *Quarterly Journal of Economics* 80: 629-649.

Roy, Donald. 1952. "Quota Restriction and Goldbricking in a Machine Shop," *American Journal of Sociology* 57:427-42.

Scharfstein, David and Jeremy Stein. 1990. "Herd Behavior and Investment." *American Economic Review* 80:465-79.

Selznick, Philip. 1949. *TVA and the Grass Roots*. Berkeley: University of California Press.

- Simon, Herbert. 1947. *Administrative Behavior*. New York: Free Press.
- _____. 1951. "A formal theory model of the employment relationship." *Econometrica* 19: 293-305.
- Spence, A. Michael. 1973. "Job market signaling." *Quarterly Journal of Economics* 87: 355-374.
- Stewart, James. 1993. "Taking the Dare." *The New Yorker*, July 26, 1993: 34-39.
- Tversky, A. and Kahneman, D. 1974. "Judgment under uncertainty: Heuristics and biases." *Science* 185: 1124-31.
- Weber, R. Rottenstreich, Y., Camerer, C. and Knez, M. 1999. "The Illusion of Leadership." Forthcoming in *Organizational Science*.
- Weingast, Barry. 1995. "The economic role of political institutions: Market-preserving federalism and economic development." *Journal of Law, Economics, and Organization* 11: 1-31.
- _____. 1997. "The political foundations of democracy and the rule of law." *American Political Science Review* 91: 245-63.
- Williamson, Oliver. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press: New York.
- _____. 1979. "Transaction-cost economics: The governance of contractual relations." *Journal of Law and Economics* 22: 3-61.
- _____. 1985. *The Economic Institutions of Capitalism*. New York: Free Press..
- _____. 1996. *The Mechanisms of Governance*. New York: Oxford University Press.
- Wilson, James Q. 1989. *Bureaucracy: What Government Agencies Do and Why They Do It*. New York: Basic Books.