Chapter 1

Alternative Roles

SUPPOSE THAT it is your responsibility to encourage more effort from an employee on a task. The task might be straightforward (say, processing tax returns or collecting mail) or more complex (say, teaching children or managing the drug rehabilitation of recalcitrant clients). Even ascertaining how much effort your employee commits may be tricky. But let's shelve those issues for the moment and ask, how do I persuade my employee to work harder?

One solution—an intuitively natural one—is to monitor the employee's performance by some objective set of standards and to reward the employee for meeting the standards. This is both an old and a strikingly modern approach to the problem. The approach has roots which extend into the very beginnings of management science—the so-called scientific management approaches of Frederick Taylor, Henri Fayol, Luther Gulick, and others. More modern branches carry the fruit of serious game theoretic approaches to the management of the firm—the approach the disciplines of political science and economics call principal-agent theory, exemplified by Bengt Holmström (1979, 1982a) and Arrow (1985).[1] But though the approach is intuitive, it is probably inadequate advice to offer to the supervisor. We will refer to these approaches as coercive, those that emphasize monitoring performance and rewarding for meeting standards (and possibly punishing for not doing so).

Instead, we may need to reach into other aspects of the theory of organizations and of management and look for alternative advice for the supervisor. The enterprise of this book is to identify what three of these roles may offer in the way of advice: training, brokering of trust, and task management. None of these roles are especially new, all date to one or another aspect of organization theory and management science. Nor are they exclusive and exhaustive, and perhaps not even entirely independent of the coercive aspects of supervision.[2] We argue, though, that they are intimately connected, and that the connection stems from the problem of organizational ambiguity.

Ambiguity is a touchstone reference for organization theorists, differing in important ways from its more familiar cousin, uncertainty. Both ambiguity and uncertainty denote states of unfamiliarity and unpredictability about outcomes. An uncertain state is one that can be resolved with additional information, whereas an ambiguous one remains unfamiliar and unpredictable even after information is acquired. An uncertain state of the world, for example, would be whether it will rain on a given day. At the end of that day, the answer is definitive. The corresponding state of the world is whether it would be good if it rained on a given day. For the farmer fearful of drought, a rainy day is unambiguously good; for the would-be picnicker, perhaps not.

This is a toy example, but it serves well to differentiate between uncertainty and ambiguity. At its most basic, ambiguity is about uncertainty of probability, and thereby serves to distinguish risk from ambiguity (Knight 1921). This is what Daniel Ellsberg shows in his thought experiment that people prefer to bet on events they know more about (1961). This is Ellsberg's paradox. Indeed, there is an aversion to ambiguity. In this guise, ambiguity is related to issues of subjective expected utility, common knowledge of rationality, and consistent alignment of beliefs (for more on subjective expected utility theory, see Camerer and Weber 1992; for a discussion about common knowledge of rationality and consistent alignment of beliefs, see Hargreaves-Heap and Varoufakis 1995). But our conceptualization of ambiguity is deeper, as our example regarding rain demonstrates. For our purposes, ambiguity is more than uncertainty about the probability of an event. It is about ambiguity over the utility of different outcomes.

The problem of ambiguity is pointed for public bureaucracies. At root, the problem of ambiguity for a public bureaucracy rises from the contradictory tasks that we, the citizens, expect bureaucracies to achieve (for the original advancement of this argument, see Wilson 1967). Bureaucracies must be equitable, and treat all inputs identically. After all, the very core idea of rational-legal authority is the idea of regularity and consistency in how the bureaucratic authority deals with the polity (Weber 1947; Finer 1941). But bureaucracies may also exercise considerable discretion, and do so because the polity believes that some cases are best handled as individuals and exceptions to rational-legal rule (Lipsky 1980).

Several other equally competing standards can be added to those of equitability and discretion. Contemporary polities prefer that their bureaucracies be honest institutions, not vehicles for the personal advancement of the bureaucrat, or, for that matter, for the advancement of privileged members of society. Likewise, an honorable goal for a bureaucracy is to be open to public scrutiny, to be more democratic in the rule of law than perhaps legislatures can be in the writing of law. But what if some sectors of a polity lack the political resources? In being fair, bureaucracies may have to cater to those without resources, to intrude politics into policy choices (Long 1949). And though it may rarely be the specific maximand of any particular bureaucrat, efficiency in the use of scarce resources may also be a bureaucratic standard (Niskanen 1975).

All of these dimensions—equitability, discretion, fairness, honesty, openness, efficiency—compete with one another, as much as they are all laudable goals. Further, it would be the height of arrogance to assert that some form of hierarchy exists among them: at times, all are candidates for the prime objective of a bureaucracy.

If these are the standards by which a bureaucrat must measure his or her performance, even one who is not self-maximizing (for wealth, perks, power, or leisure) is caught on the horns of ambiguity.

We see three basic questions confronting the subordinate in a bureaucracy:

- What are the bounds of what I can or cannot do?
- What should I do first? What is most important?
- What happens if I exercise discretion over either of these first two choices?

The first is not merely a question of aversion to sanctions or limits of authority. As the classic work by Daniel Katz and Robert Kahn laid out, *role ambiguity* was defined "as lack of information regarding supervisory evaluation of one's work, about opportunities for advancement, scope of responsibility, and expectations of role senders" (1978, 190). Elements of the problem of role ambiguity are as old as any concept in the social sciences—Émile Durkheim's idea of *anomie*, for example, is derived from a sense of normlessness in the allocation of rewards and opportunities (1893/1997). Consider one of our core examples in this volume: social workers are swamped with problems of role ambiguity, which stem from the contradictory expectations that the polity heaps on performance. Although we may have warm fuzzy images about social workers, and the nurturing aspects of their jobs (for example, extending welfare benefits, providing counseling), we also have highly punitive ones about other aspects of their work (for example, removing children from homes, assignment to mental health facilities).

In a recent book, Bryan Jones (2001) notes that ambiguity in public life can create situations in which not only the problems but also the potential solutions are ambiguous. This formulation extends from the famous work by Michael Cohen, James March, and Johan Olsen (1972) on the inflammatorily named garbage can model of organizational choice.

Jones's key elaboration is that knowing what to do may reference the tasks of advocacy of particular remedies, implementation of those remedies, or even identification of the problems.

In their article, James Breaugh and Joseph Colihan (1994) test the construct validity of the role ambiguity concept. Part of the difficulty with the concept is that it remains to be established whether role ambiguity is inherently noxious (see Kahn et al. 1964; Kagan 1972; Sorrentino and Short 1986), or is noxious because it interferes with functional performance (Sawyer 1992). Breaugh and Colihan identify three specific types of job ambiguity: employee uncertainty about the standards a supervisor uses to determine whether job performance is satisfactory (performance evaluation ambiguity), uncertainty about the methods or procedures one should use in work (work method ambiguity), and uncertainty about when to perform which tasks (temporal role ambiguity). Using confirmatory factor methods, Breaugh and Colihan establish that the three types of job ambiguity are indeed separable, correlated not only with one another, but strongly also with job satisfaction and satisfaction with the employer. Likewise, Linda King and Daniel King (1990) also see multiple facets to role ambiguity: what is required, how responsibilities are to be met, uncertainty about whose expectations must be met, and with what consequences. King and King are, however, much less sanguine than Breaugh and Colihan about the construct validity of the role ambiguity measures, and conclude that there is considerable ambiguity about role ambiguity.

On the other side—the effects of role ambiguity on job satisfaction compared with other potentially noxious qualities—there is equally strong evidence. John French and his colleagues (1982) argue that job strain results principally from a misfit between the workplace environment and the person's attributes, instead of responsibility or work. In a reestimation of the same data using more sophisticated confirmatory factor analytic techniques, Jeffrey Edwards and Van Harrison (1993) find that, although one can separate job strain into multiple dimensions (complexity, role ambiguity, responsibility for persons, and workload), only role ambiguity consistently provides a strong explanation for job dissatisfaction.

One possible mechanism for the noxiousness of role ambiguity is that it may heighten sensitivity to potential threats. In a study of the relationship between anxiety and ambiguity, Colin MacLeod and Ilan Cohen (1993) find, in a set of experiments on undergraduates, that anxious subjects in a state of role ambiguity are more likely to read text (and, presumptively, other messages) as threatening than those who are not anxious about role ambiguity. A second postulated mechanism (Webster and Kruglanski 1994) is that some subjects may have a greater need for cognitive closure about tasks, much as some people experience a greater need for cognition (Cacioppo and Petty 1982).

The second question—what should I do first? what is most important?—arises because few bureaucracies have the resources to accomplish all of the tasks within their purview. Prioritizing what has to be done when—and, as a necessary consequence, deciding what will not be done ever—is a fundamental ambiguity for bureaucrats. We devote several chapters to the problem of task allocation, which is one way to resolve task ambiguity.

The final question—what happens if I exercise discretion over either of these first two choices?—concerns political accountability. Where the Weberian ideal-type bureaucracy required clear lines of political accountability, few real bureaucracies feature single threaded lines of command. Returning to our social worker example, the subordinate social worker has a supervisor to whom he or she reports, but may also come into conflict with subordinates or supervisors in other divisions, with the director of social services, with local political officials, with reporters for the local media, and with citizens directly. The lines of political accountability become even more tangled when one accounts for the competing political interests of these many sources.[3]

Organizational ambiguity describes the conditions that arise in trying to answer the two fundamental questions every subordinate employee will have about the supervisor's expectations for performance:

- What do I do?
- What happens to me if I do it?

Organizations swim in ambiguity, and public organizations especially so.

We ask bureaucracies to do too many things, with too few resources, and where even those things may represent contradictory demands. Consider the public school math teacher. She is supposed to accomplish the core mission: to educate the pupils in a specific subject. She is probably underpaid, and probably severely so. She probably has too large a class. She may have a chalkboard. She is supposed to keep order in the classroom, which may be all the more difficult given the typically large class size. She should monitor the behavior of her students, she needs to attend to the needs of those who are having difficulty with their homework, needs to keep the attention of those who are working at the pace of the class, and even needs to keep the super-high achievers still motivated.

The problem is that we have little agreement on just what she is supposed to cover. As all of us who have been through public schools may (perhaps painfully) recall, a typical progression in math classes is from basic skills, through algebra, analytic geometry, advanced algebra, trigonometry, pre-calculus, calculus, and perhaps beyond. But some schools may, perhaps rightfully, argue that our teacher needs to cover calculus just before college, and insert another topic (discrete math and statistics are common). Other schools (and parents) may expect teachers to combine some of these topics. Still others may prefer calculus to come earlier in the academic sequence. The question of the best way to teach is ambiguous because it does not have a definitive solution.

That said, it does not stop us from imposing at least some standards on performance—that of the pupils, the teachers, and the schools. The No Child Left Behind Act calls for mandatory end-of-year testing of children's performance in public schools, ostensibly to permit students to relocate away from failing schools. Similar policies are common in the states. This is an exemplar of the coercive approach to public management: insist on the output of the public school (and, more directly, the teacher) be measured by some objective set of standards, and punish (by permitting the students to relocate to other schools) schools that do not achieve.

A funny thing happens. Schools, and teachers, are quite familiar with the year-end tests and the consequences of inadequate performance. The teachers may explicitly adopt a policy of teaching to the test, that is, where the sole focus of the year's instruction is the test itself. Similar policies may emerge in the nation's AP exams. This is one answer to the problem of measuring performance and an instantiation of achievement.

It is not, however, necessarily the right answer. What about those problem students, who are found in every school in the country? When the end-of-year tests come around, teachers may choose to "cheat" on them by encouraging slow learners and distracting students from attending school that day. Some students may have learning difficulties, such as attention deficit disorder, dyslexia, or dysgraphia. How much should we adjust the test, either in scoring it or in how students take it, to accommodate such students? Is teaching to test an appropriate way to encourage more ambiguous forms of learning, such as fostering a love of learning itself?

Public bureaucracies face exceptional ambiguity, in that supervisors and employees must make choices about unattainable goals that often contradict each other and for which we underequip them.

Our first book on performance in public bureaucracies, *Working, Shirking, and Sabotage* (Brehm and Gates 1997), claimed that to understand which bureaucrats put positive, zero, or negative effort at their jobs, one should really look more to the bureaucrats than to their supervisors. And so, as many readers have commented, this raises the question as to why we even have supervisors. A justly appropriate question, but not one we originally considered. Our purpose in this volume is to suggest that supervisors do meet very functional, very desirable purposes within public organizations, just not those that fall squarely within coercive conceptions of the organization.

We will retain our dyadic approach to thinking about supervision. Our focus will be on individual bureaucrats, but in thinking about the alternative conceptions of the supervisor. We believe that such an inside the box examination is vastly superior to the highly aggregated approach.

In the remainder of this chapter, we first provide an outline of the ancien régime, the coercive approach. One should, after all, really understand what one is rejecting before rejecting it. We then review the three alternative approaches to management—training, trust brokering, and task management—and explain some alternative (and why we choose not to cover them here). Finally, we outline the broader plan of the book.

The Coercive Approach

Principal-agent models focus on an individual (the principal) who contracts work from an agent. The agent's motivations and ability (or type) may be unknown, as may be his or her effort on the job. The former—unknown type—is classified as a problem of *adverse selection.* The term comes from the insurance literature, and we credit Kenneth Arrow (1985) for articulating it. The latter problem—unknown action—is one of *moral hazard,* another term from the insurance literature.

In the insurance industry, one problem for the company will be that some individuals may buy insurance only because they expect to make a claim. The desperately ill patient is a classic example. This is adverse selection.

In an organizational setting, adverse selection refers to a supervisor's inability to gauge the subordinate's preferences and abilities. How might the principal resolve this problem? Perhaps by verifying the agent's type through intensive interviews in screening candidates, or by probationary periods for hires, or by expectations of certain qualifications for the job, similar to the mandatory blood tests before securing life insurance.

The moral hazard problem also affects the insurance industry. Here, instead of purchasing insurance because one expects to need it, the purchaser changes his or her behavior. In an organizational setting, the action of subordinates focuses on degree of effort (or propensity for shirking) and latent opportunism. How can one resolve the problem? One might monitor the subordinate's behavior, or change the incentive structure (by increasing penalties for bad behavior or increasing rewards for good behavior), or even the entire system in favor of incentive-compatible regulations, just as the insurance company might reduce rates for drivers with good records or dispense rewards for those who can prove they do not smoke.

Sometimes the problem may be particularly pointed, as when there is a team production condition (Alchian and Demsetz 1972), where all the principal sees is the total output from a team of agents. At this, the incentive to free ride on the contributions of others becomes an additionally strong component of the moral hazard problem. Bengt Holmström (1982a) extends this analysis by demonstrating that there is no budget-balancing revenue-sharing scheme that can effectively induce all members of a team to work and for none to shirk. In other words, ultimately, a principal is stuck with the moral hazard problem.

The intuitive way to rectify both the adverse selection and moral hazard problems might be to improve the selection of employees, monitor their output on the job, and pay them according to that output on the job. An early advocate of scientific management, Frederick Taylor, argued for just that approach, which was grounded in four principles: accumulating all the knowledge about the job, scientifically selecting the workman, offer incentives for good performance, and reconsidering the division of work itself (see Taylor 1911). These themes of coercion appear again toward the end of the twentieth century in work related to organizational oversight in the public sector (see, for example, McCubbins and Schwartz 1984; Fiorina 1986; McCubbins, Noll, and Weingast 1989). Roderick Kiewiet and Mathew McCubbins feature a set of coercive powers available to supervisors: hiring and firing, design of compensation (contracts), monitoring and institutional checks (1991). These have a broad intuitive appeal.

Indeed, the approach is intuitive, but it is also incomplete. Chester Barnard, an executive at New Jersey Bell Telephone, suggested that organizations function more as cooperative ventures rather than coercive ones. Indeed, he wrote of the "fiction of the superior authority" (1938/1968, 166). Although his book, *Functions of the Executive,* serves as something of a model for this volume, Barnard is too harsh in his assessment of the fiction of coercive authority. Coercive authority can work: people do lose their jobs due to underperformance, are paid piece-rate, and are selected after careful evaluations of their capacities. We argue, however, that such coercive roles, as emphasized first by Taylor and later by scholars of the economics of the firm, are misplaced in studies of public bureaucracy.

The coercive model has limited applicability to private sector firms, but even more profound ones in regard to public sector organizations. First, public institutions are not governed by the laws of supply and demand. Except in cases where they compete to provide services, bureaucracies tend to operate as monopolies prone to oversupply policy output. There are revenues and budgets, but no profits (and with implications for principal-agent models, no residuals). Such limitations call into question whether principal-agent models apply to the public sector. Second, civil service laws significantly constrain public sector executives from using coercion in any form. A supervisor has, unlike his or her private sector counterparts, limited ability to redistribute pecuniary rewards to reward or sanction subordinates. Moreover, the ultimate enforcement—firing someone—is extraordinarily difficult and typically possible only in cases of gross negligence or misconduct. Third, bureaucratic policy output is notoriously difficult to monitor. Public sector policy output typically does not lend itself to piece-rate production. Ultimately, public sector organizations are inherently infused with ambiguity. There is almost never a bottom line, but rather many overlapping and competing goals. All organizations, even small private sector firms, face problems of ambiguity, but public sector organizations must overcome more fundamental issues of ambiguity—ambiguities in terms of goals, design, and organization. Clearly, the coercive role of a public sector manager is even more limited than that of a supervisor in the private sector. What then, are the functions of the public executive?

Alternative Functions

Let us look at three alternatives to the coercive roles of management: training, cultivating trust, and task management.

By training, we mean two distinct things. One is the cultivation of prosocial preferences: to learn to want to do what the principal wants. This quite possibly happens in some organizations, though the record is rather mixed. Another meaning is the clarification of the boundaries of acceptable and unacceptable behavior. Every organization has rules, whether formally denoted or not. As Charles Perrow (1986) argues, organizations without formal rules require a fair amount of exploration by the organization members, and usually entail some significant pain to discover. We explore both aspects of training.

By cultivating trust, we mean one thing: to what extent does the supervisor gain the trust of his or her subordinates, and by what mechanism? The fuzzy part of the argument might be with the mechanism. Sometimes, subordinates express trust in their supervisor because of a belief about his or her intentions, or one about the supervisor's character or capacities. We argue that supervisors can gain the trust of their subordinates by providing cover for them: acting as interference to prevent others from interacting directly with the subordinate (such as members of the media, other supervisors, or other political superordinates). This is a specific trust owing to a specific content.

By task management, we refer to the capacities that supervisors may have at their disposal to persuade subordinates to work on some subset of tasks at the omission of effort on others. The supervisor may prefer that the subordinate allocate effort on tasks A and B instead of tasks C and D, whereas the subordinate may prefer to work on tasks B and D. In this scenario, the supervisor's leverage is of effort at B.

We believe that a full understanding of just what we mean by training, trust brokering, and task management, an overview of the chapters to follow is warranted. But, first, let us give a preliminary sketch of the idea that motivates our interest in these specific functions, as well as why we believe they are intimately related.

Sometimes, taking a job in the public sector entails the possibility that a great number of people and institutions will be watching what you do and how well you do it. Sometimes the rules for what you can and cannot do are explicit, but other times they require clarification. Sometimes they contradict one another. How does one respond?

One response, if you can get away with it, might be do to as little as possible. This is what we refer to as shirking, or leisure-shirking (Brehm and Gates 1997). But another response, especially for public sector employees motivated by functional preferences, would be to contribute to the purposes of the organization, perhaps even beyond how the organization tangibly rewards you. Let's call this working. There is also a third, less commonly observed response, which would be to consciously undermine those rules through diligent efforts, which we have referred to as sabotage (Brehm and Gates 1997).

How do you know which rules (formal or informal) are the most important to follow and when? The first function of the public supervisor that we explore is the role of training, understood here as the clarification of boundaries. Of course, other parties, such as one's co-workers, may be even more important in clarifying the rules. If the supervisor takes the lead role in training the subordinates, then folk cultures may be more under control and the norms for working instead of shirking can take charge.

There is a flip side to the supervisor taking such a strong role in training her subordinates: she has to stand up for them against all other parties who may be interested in their behavior. Other supervisors may put demands for diversions of the subordinate's time or just disagree with how this particular supervisor sets her priorities for performance. The local media hardly ever report about the successes of bureaucracies, most of which are routine and mundane and ordinary. But bureaucratic failures illustrate the alleged heartlessness, cruelty, and arbitrariness of those bureaucracies and make for great headlines. Local government officials may also become interested before or after such coverage. There are many people and entities out there who can make your performance a cause for personal grief.

We call this second role trust brokering. Supervisors who stand up for their employees by providing political cover against those other interested parties earn their employees' trust. This is hard work and may be even politically costly for the supervisor. But brokering trust earns something, namely, that those subordinates who feel they have the discretion to do what is right, do it, and have the full faith and confidence of their supervisor, and accord their supervisor leeway.

Leeway is essential. As we have written and will further explain in this volume, bureaucracies generally do many different things. Of course, some perform a highly circumscribed set of tasks, such as the Social Security Administration or the U.S. Postal Service. James Wilson referred to these as production agencies (1989). Most bureaucracies have more to do than time to do it.

Supervisors must therefore choose among these tasks, setting priorities, perhaps encouraging subordinates to perform less interesting or even onerous or unpleasant tasks. Given the limited flexibility of the public sector, compared to the private sector, such allocations are extremely important for the public executive. This third function is task management.

The rough summary of our argument is this: because of the fundamental background of ambiguity behind public bureaucracies, there is a key function that supervisors must perform in training their subordinates about what is allowed and what is not; in exchange, the subordinates only trust their supervisors when that supervisor shields the subordinates from intrusions by others; in doing so, the supervisor gains more leeway to assign tasks.

Plan of the Book

We develop this argument by taking several small steps, applying a mix of methods, models, and data.

Chapter 2 looks to empirical data on the effectiveness of training in real bureaucracies, here, drawing on our own data on social workers in North Carolina and three waves of extensive surveys of employees in the federal bureaucracy.

Chapter 3 explores the likely futility of training as a cultivation of pro-social preferences. Our tool is a computational model that closely parallels the structure of a basic principal-agency relationship.

In chapter 4, we develop a cooperative game theoretic model of the task allocation problem. We prove the existence of a cooperative solution, called a core, within a budget that depends solely on exchanges of perks for tasks.

In chapter 5, we analyze the task allocation problem in practice, in a study of task allocation within the police forces of three separate cities.

In chapter 6, we articulate a very simple (noncooperative) game theoretic model of trust, based on an extension of the trust-honor game by David Kreps (1990), and again explicating this model in the context of the social worker data.

In chapter 7, we explore the same problem—that of the cultivation of trust—by examining the behavior of federal bureaucrats.

Chapter 8 unifies the three functions and argues that the unique problems of ambiguity for subordinates in public bureaucracies create conditions under which training clarifies boundaries, trust results from supervisors standing up for subordinates, and trusted supervisors have greater latitude in assigning tasks. We conclude with speculation about the generalizability of our findings to the broader problems of leadership within organization, and the organization's performance as an aggregate.

This volume draws on three sources of data from different types of bureaucracies: three surveys of federal government workers conducted by the Office of Personnel Management (OPM) in 1998, 1999, and 2000; participant observation data on police officers (Ostrom, Parks, and Whittaker 1982); and our own surveys of social workers in North Carolina. We use three different theoretical modeling techniques: simulations or agent-based modeling; noncooperative games; and a cooperative game. We also draw on a variety of statistical models, including first principal component scale construction, mixture models, and three-stage least squares.