1 Information, Uncertainty, Structure, and Function in Organizational Sociology

Rationality

Rationality necessarily involves an analysis of the future, because the consequences that give purpose to acts are necessarily in the future. Thus all rationality is based on predictions of one kind or another, not on knowledge. Assuming that actors are perfectly rational, of course, implies that they are certain what the future holds, that in all relevant respects our notions about the future constitute knowledge. The assumptions of neoclassical economics are that various financial quantities (e.g., the interest rate and the savings and investment rates if distinct) summarize all the relevant information about the future that one needs in order to make economic decisions, and that everyone else can have the same information that an individual actor uses to work out a strategy (the investor, for example, presumably compares a concrete opportunity [e.g., to invest] with the relevant financial quantities [e.g., the interest rate]; the assumption is that anyone else with resources has access to the same information about the future relevant to that investment as the investor does). The assumption of neo-Keynesian economics (and apparently of Keynes himself) was that the idea that investment can operate with only financial and universally available knowledge assumes a kind of knowledge about the future that human actors do not and cannot have (see Weintraub 1979).

Two main traditions of organizational sociology start from the Keynesian or neo-Keynesian assumption that the future is uncertain. The older one has as its stellar figures Herbert Simon and James G. March (e.g., March and Simon 1958). Its main thrust has been to try to describe what people acting in organizational contexts (and to some degree acting in organizational roles in the light of organizational purposes) actually do to deal with the imperfection of their knowledge or predictions of the

future, what decision procedures they in fact use to fill in what neoclassical economists require them to know.

The other tradition, whose stellar figure is perhaps Oliver Williamson (e.g., Williamson 1985), looks at what organizational devices might be good for dealing with various sorts of uncertainty about the future (e.g., small numbers bargaining, which involves uncertainty about the future behavior of a partner who is not replaceable in the market—or is replaceable only at a higher cost). The organizational devices solve specific types of departure from perfect markets. The actual distribution of such organizational devices, then, is to be predicted from the distribution of the problems they are good for solving. The Williamson alternative is, in short, microeconomic functionalism. The approach to the microeconomic ideal of perfect efficiency in spite of uncertainty is the survival criterion for the practice in a functional argument.

To put it another way, the Simon-March tradition is concerned with the *causal explanation* of departures from rationality. The Williamson tradition is concerned with the *functional explanation* of organizational structures, such as hierarchies. The Simon-March objective is to explain why organizations use less than rational methods. The Williamson approach is to explain how organizations can come near rationality, if not quite there, by adopting organizational devices that deal with uncertainty. The clue to structures needing such explanation in the Williamson tradition is that they occur in the context of market behavior but depart in significant ways from the structure of decision making known to be optimum in a perfect market.

I propose in this book to follow a slightly different strategy from both of these, though the difference involves a minor shift from the received tradition. I propose that information about the uncertain future becomes progressively available *in distinct social locations*, depending on what sort of uncertainty it is. What resolves the uncertainty of particular actors, then, is the *earliest available information* that will show what direction the actor ought to be going because of the way the future of the world is, evidently, turning out. Structures that depart from idealized markets (none of the organizational structures we explain in this book work like ideal markets do) are then to be explained functionally, by the growth of the organization toward those locations where information for resolving uncertainty is chiefly located. That information then has to be processed quickly, both to adapt the previous tentative strategy and to work out the tentative strategy for the next period. The core of structure of organizations, then, is information processing, and the core information to be processed is the earliest information that indicates what sort of world (i.e., future) the decisions are being made for.

The explanation is, however, causal as well as functional. The organization controls neither where the uncertainties it is confronted with arise, nor exactly what sort of information is crucial. Those are determined by what the organization is trying to do and varies from one part of the organization to another, depending on what that part is trying to do. That is because different things one is trying to do have different bottlenecks, different main loci of uncertainty about the future, and consequently different information that will bring success.

It is, then, the factual distribution of information relevant to crucial uncertainties (crucial in particular, in economic organizations, for the net value of the object or service produced by the organization) that determines the actual structure of the problem confronting the functional structure. Structures of organizations, and of parts of organizations, vary according to the sorts of uncertainties they confront, and so according to what sources of information they depend on and to how that information is best got to the decision-making units. The board of directors is organized as a committee, the assembly line as a hierarchy, because the financial and trustworthiness uncertainties faced by the board are different from the labor cost uncertainties faced by the assembly line.

The crucial thing for an organization from this point of view is to be where the news breaks, whenever it breaks. Information is "news" for the organization when it is a first appearance of some sign of how the future is going to be, in a respect crucial for the organization. The argument then is that news about one sort of contingency facing an organization (e.g., fluctuations in the quality of raw materials) is likely to be found in a different place than news about another sort of contingency (e.g., the risks of a hostile takeover). One therefore needs a type of structure to gather and process information from raw materials sources to manufacturing different from that needed to gather and process information about suspicious movements of stock sales.

In both cases, however, the organization does not know what kind of world it is making decisions in, because it cannot predict the future very well—not well enough to stay afloat. It needs the news, and it needs to respond to the news as it breaks, not after the news gets to all the participants in the market in the form of price changes. It is in fact such reactions to news that inject information about what is happening to corporate futures into the price system.

Much of this book will apply the above approach to economic organizations, partly because the classical literature on economic organizations has been sensitive to risks and to information processing and partly because economic organizations is what I know most about. However,

the language above may suggest more restriction of the problems of uncertainty and information processing to economic organizations than is appropriate. Certainly a university scientist, in order to establish that brief moment of monopoly reflected in his or her name as author on the first publication of a new finding, is highly uncertain beforehand about where new knowledge is to be found (see, for example, the discussion in Watson's *The Double Helix* [1968] 1981) of getting the x-ray diffraction results from Rosalind Franklin, which limited the possible structures of DNA) and how likely it is that someone else will get there first. Such a scientist (and therefore the university) must be very attentive to the news of scientific results related to his or her line of research. That uncertainty of research is compounded as deans and department heads bet on who is going to make the correct bets about where new knowledge is to be found, without themselves being able to make those bets as well as the people they are hiring (for example, see the evidence of wrong personnel bets by Watson and Crick's administrative head, Sir Lawrence Bragg, in Watson [1968] 1981, 37–40, 60–62).

Universities also tend to grow toward sources of news that tell them as soon as possible what the future of science and scholarship is going to look like, and to organize their "administration" so that decisions on research and hiring are taken primarily by people who have efficient channels to that news. We will treat a few problems in university administration in Chapter 9 to show that the approach taken here extends beyond economic organizations.

Uncertainty

The basic idea here, then, is that uncertainty is not an indelible characteristic of a particular decision. Instead it changes over the course of the decision. That is why people take decisions tentatively. They enter on the first branch that might lead to an investment (e.g., a "conceptual study" in engineering) in order to get part of the news that will reduce uncertainty (see Marschak, Blennan, and Summers 1967, 49–121, on the process of decision to follow or not follow a line of development in designing jet engines; see also Crecine 1969).

In general, the further a commitment to a line of action extends into the future, the larger the proportion of all the information entering into the final decision will be that is news coming in over the course of implementation of the decision. People do not decide to drill exploratory wells until after geological studies have shown a promising formation; they do not drill the first production wells until exploration shows that the find is

"commercial" (i.e., that its rate of flow, costs of development, transportation costs, etc., combined with anticipated oil and gas prices over the course of the wells' lifetime, show "adequate" anticipated profitability); and they do not develop the whole field until the first production wells come in as anticipated (see Stinchcombe 1985e, 41–65, for a more detailed analysis of the uncertainties in oil field development). That is, they are not nearly as blind to the nature of the future at the time a major investment is made as they were when they first set out with sonic equipment to explore the shape of the rocks a couple of miles down from a given pasture.

But of course, at the very end OPEC may fall apart and the field turn out not to be commercial after all (Herskedal and Kristiansen 1987). At this point, after one has made tentative decisions to go a bit further, waiting for the final commitment until as much news as possible comes in, then one has to estimate the probability distribution of oil prices and of true rather than estimated construction costs, and "take risks." Uncertainty is reduced through news; then finally the residual uncertainty is transformed into risk, and people make their bets. Or perhaps better, people make small investments and build a small structure to collect relevant news; if the news is good, they make bigger investments and develop a larger structure to collect relevant news; and so on. Uncertainty is transformed piecewise into risk, with a large part of the risk at first being a guess concerning the value of the news that a news-collecting structure will bring in.

This orientation to news must partly explain the fact observed by Mintzberg (1973) that top business managers prefer their information "live," from conversation, telephone calls, inspections, problems brought in by subordinates. Top managers are primarily responsible for dealing with uncertainties, because they are nearest in the hierarchy to profit takers who ultimately take the financial risks. So the parts of the total communications structure they are likely to be most attentive to are the parts that follow the fast-breaking news, the news that reduces uncertainty. It will be important for them to use fast communications channels, voice and telephone, rather than waiting for a memo to be drafted, revised, and forwarded with appropriate signatures up the line.

Uncertainty About What?

The primary task of this book is not to follow the drama of taking risks in the light of the last possible reduction in uncertainty. Rather, it is to analyze the structure of organizations as determined by their growth

toward sources of news, news about the uncertainties that most affect their outcomes. The basic notion is that those uncertainties are distinctive in different parts of the organization, depending on their distinctive tasks and environments. A way to conceive of this is as follows.

By analogy with the findings of linear programming methods for optimizing complex decisions, we can imagine that any particular process has from two to a few "constraints" on its success. At any particular instant, there are likely to be only one or two variable factors that will improve short-run success (such as persuading all the workers to come in Monday morning on time), but those may change from instant to instant over quite short time periods. Over longer periods there may be several constraints which may be overcome or removed, at greater or lesser expense. We can for convenience call these the "bottleneck" factors in a given type of work with a given set of historically derived assets. Many of the things that are essential for the process to go on are not constraints at a given time (ordinarily, for example, air for workers to breathe is not a constraint; at times in underwater construction, however, it is the crucial constraint).

The basic generalization is that only uncertainties about the constraints on success matter, and so only information about those constraints on success is crucial to the operation. Because the constraints are likely to change over time, the uncertainties that are relevant and the news that resolves those uncertainties are likely to change.

It may often be, however, that the changing information is structurally located in the same place, so that the same structures will carry it back to the organization. For example, the information relevant to figuring out the structure of DNA changed as the alternative hypotheses under consideration changed with the gusts and flurries of facts and ideas (Watson 1968). But most of the information relevant to *any* of the alternatives was located among physicists and physical chemists interested in biological problems, especially the "phage group" organized around Delbrück and Luria, the crystallographers working on proteins organized around Wilkins and Franklin, and the physical chemists organized around Linus Pauling. The informal network bringing in the news did not have to be rebuilt with each shift in hypotheses, because basically the same sources of news had the facts about each of the alternatives.

Similarly, as we will examine in more detail in Chapter 5, information about needed improvements in an innovation that will make it more marketable comes from those who are in contact with the innovation users. In the early days of computers, given the structure of maintenance agreements and software development at the time, these were likely to be

the maintenance people or the "client representatives" of the vendor of a computer, the folks who would be in contact with the computer center of the vendor. So it was essential that information from the client contact people get back to the software and hardware developers in the vendor organizations, whether the design problem was with the software, the mainframe computer, or the peripherals, whether it was in the early days of the computer model or the last days before it was replaced by new models, and so on.

It is such repetitive use of information from a given source of news, as problems and constraints on success change and new uncertainties therefore affect decisions, that tends to cause dense and fast information collection and information processing in organizations to be oriented to particular sorts of uncertainty and particular sources of information. It therefore follows that different organizations, having different main constraints, hence uncertainties, and hence news sources, will need to grow different types of structures of information collection and processing. The basic functional postulate of the argument of this book is that such need produces a tendency to grow such specialized structures for processing information about the relevant uncertainties that differ in different organizations.

But the same logic applies to the uncertainties of different tasks in different parts of the organization. Different uncertainties matter for different products, for reaching different clients successfully, for arranging the flow of borrowing and repayment so as to give stockholders their optimum risks and rates of return (and so the maximum value of their stocks) in different states of the financial markets, for maximizing university prestige in molecular biology versus other parts of physics or other parts of biology, and so on. We will therefore expect that the structures for information collection about crucial uncertainties (and especially that part of information processing that involves rapid use of fast-breaking news) will in general tend to be different in different parts of the organization.

Information

Up to this point we have emphasized the temporal features of that sort of information that reduces uncertainty, that what is uncertain at one time becomes predictable (or becomes calculable risk) at another time because new information (news) comes in. But many other features of information besides its recency make it functional for reducing uncertainty. Consequently, many other features of an information-

collection and -processing apparatus—besides being located where the news about the relevant uncertainties is breaking and getting that information fast to the relevant decision makers—determine the degree to which they "serve the function" of controlling uncertainty, turning it into calculable risk.

We will discuss these secondary features under the rubrics of: (1) units of analysis; (2) noise reduction; (3) level of temporal abstraction; (4a) trustworthiness and error estimation and (4b) agency problems in information trustworthiness; and (5) expertise of analyzers.

1. *Units of analysis.* The kinds of causal unities about which an organization needs coherent information are those that make several activities subject to the same constraints: then units made up of those activities have the same conditions of success. For example, in studying production delays in a steel mill in South America, I found that the hot rolling section of a tube mill collected and analyzed delay information but that the cold processing parts did not and their delays were not treated as part of the problem of the efficiency of the plant. The reason was that on hot rolling operations there is a high degree of technical interdependence of the different operations, and they have to be done in tight sequence because otherwise the steel will cool down so much that it cannot be worked. That meant that the success of any part of hot rolling depended on delays in every other part, so any constraints that caused delays in one suboperation were a constraint on all the others (Stinchcombe 1974; Stinchcombe and Harris 1969). It was very uncertain whether at any particular moment the line would run, because many interdependent difficult operations all had to run at once for the line to run.

Similar inputs (e.g., preventive maintenance for the whole hot line) would tend to remove all the constraints on operation, so information about maintenance problems of any part of the hot line would be useful in forming a remedial policy. Incentive systems that encouraged one shift not to leave steel hot, ready to be rolled, depressed productivity of the whole line on the following shift. In short, the technical unity of the line meant that it was subject to common constraints, that there were common remedies for those constraints, and consequently that common information would be valuable in increasing productivity of each of the parts.

For another example, we treat below in Chapter 4 the question of what Chandler means when he says that different lines of merchandise will require separate administrative divisions. One of the examples he gives is Pittsburgh Plate Glass. The firm added several lines of merchandise that tend to be used with plate glass, such as paints, brushes, putty,

and the like, which have no technical unity with plate glass *in production;* they do not need to be heated and formed while hot or shipped in special ways because of their fragility. But the information about what is selling to finishing contractors in lumber yards and hardware stores around the country is located in the same place for glass as for putty, and for both of these as for paints and brushes. Consequently, the lack of technical unity in manufacturing did not require divisionalization of Pittsburgh Plate Glass.

When the firm started to sell some products to other manufacturers on a large scale, however, those products and their manufacturing and marketing staffs needed to be divisionalized because they had to respond to a completely different structure of uncertainties and information-processing requirements. The news about what was selling in lumber yards was not useful in forming an inventory policy about such industrial goods. The market constraints on plate glass and paint brushes, then, tended to be similar, and to be solvable with information from the same structure of contacts with lumber yards. The market constraints on industrial chemicals were not similar, and so they had to be placed in a different division.

For Sears, analyzed also in Chapter 4, the statistical system for inventory control in stores requires each distinct commodity to be treated separately, since all the thousands of commodities that Sears sells in its department stores are bought separately, reordered separately, inventoried according to the expected rate of sale and the reorder interval, and so have to be treated as separate units of analysis. The "unit control" system developed in Sears retail operations provided such a set of statistics on the sale and inventory of separate products.

There are two general points to be made about this. First, one can support a functional argument that an information system has its origin in a need to respond to a given sort of uncertainty, *if* it has the units subject to that sort of uncertainty as its units of data collection and analysis. Conversely, one can assume the functional argument and infer from the units of the information system what the uncertainties are. We infer from the units of a preventive maintenance system that the hot rolling section of a tube mill has as a principal constraint the maintenance of all its parts, while the cold processing operations are not subject to the same constraints and so do not need information about the same uncertainties.

2. *Noise Reduction.* The odds at the race track and the average price of a stock over a week or so are good estimates of the long-run relative merits of horses or of the expected earnings and risks of a corporation. But almost all the bets of particular people, given the odds, and almost all the minute-

to-minute or hour-to-hour fluctuations of stock prices are error noise around that good estimate. Race track touts and investment analysts cannot ordinarily beat, respectively, the odds or the market, so their advice about what horse to back or what stock to buy is random noise. What one wants in dealing with uncertainty is a system of fast analysis (as the track itself has for odds, for example) that will eliminate the noise and error and provide only the valid information about the constraint. IBM, in pricing a personal computer, does not want to know that a particular retailer always undercuts the official price by 3 percent unless it can learn how the retailer can afford to do that and then use that information to its own profit. But the company does need to know if all retailers are cutting the price by 3 percent rather than 10 percent, because 3 percent means it is meeting the competition in its pricing policy, 10 percent that it is not.

A central kind of noise is irrelevant detail in a report. A twenty-page output from the accounting department for department or project administrators has to be reduced, often by pencil calculations (usually by administrative assistants or the equivalent), to about three or four numbers that tell where they stand. In analyzing an information-processing system, then, one wants to know the four numbers that the best administrators of departments or projects have calculated for them, rather than what all is included in the twenty-page biweekly output that is stacked in a closet (after those four numbers are abstracted) until the auditors say it can be thrown away. Such detail is comparable to knowing how much each bettor has bet on each horse rather than knowing the summary (the track odds).

Another kind of noise is, of course, systematic error. Deans complain, for example, that all professors they talk to earn less than average for people of their rank and distinction. The fact about which they complain comes about because professors have reasons to select the set of people or institutions to average so as to make the average as high as possible without seeming absolutely ridiculous (see 4b below on agency). We will deal with the general problem of systematic error in the section below on trustworthiness, and reserve the term *noise reduction* for reduction of random error.

The basic device for reducing random error is to take a mean (or median or percentage); the larger the number of observations over which the mean is taken, the smaller the residual random noise. But a larger sample of fast-breaking news is always late. Consequently, the dilemma is to produce the most recent mean possible, compatible with its being sufficiently noise free to indicate what is happening to the crucial uncertainty, to indicate how the world is developing. If an information system

shows signs of this tension between noise reduction and newsworthiness, for example by using a four-week projection of demand updated each week (as General Motors developed—see Chapter 4 below), this shows that it is oriented to getting noise-free information about that uncertainty that is recent enough to adapt to a changing and uncertain world.

The alternative strategy for dealing with noise occurs when each case has to be dealt with in the way it in fact turned out. In the textile industry, the mean strength of a thread being spun or woven does not tell when it will break; the weaver-loom system has to be set up, then, to deal with each break in the thread after it happens, and taking the mean time to failure only helps in buying thread, not in running looms (Blauner 1964, 59). The machine has to be stopped and the break tied up when it in fact happens; one cannot set up the work on the shop floor based only on the fact that an average stretch of thread in a spinning machine will not break.

Skills in the division of labor are primarily for dealing with variation in the situation: in the raw materials, in the design requirements, in the previous errors (for instance, previously built walls not being quite plumb), in the weather, and so on. The reason one needs the skill of a carpenter is that one gets the news that the wall is out of plumb when one is hanging the door, and so one needs to deal with it at that point. In Chapter 2, we deal with the skill composition of the labor force as a response to substantive noisiness of the work process, the amount of random variation that has to be dealt with rather than averaged out, drawing heavily on the work of Charles Sabel (1982).

3. *Level of Temporal Abstraction.* Different parts of an organization deal with uncertainties that develop over different time spans, and different organizations have crucial uncertainties that vary over different time spans. The abstractions built into the information-collection and information-processing structures have to be adapted to the time spans over which adaptation takes place.

For many manufacturers, much of their capital investment is in buildings that are reasonably easily turned to other purposes, with the old machines moved out and new machines moved in. But oil and chemicals firms build special-purpose capital equipment (often roughly 5 percent only is for the building, 95 percent being for the processing installations); such special equipment has to pay for its entire cost by a long time doing the single thing it is built for. Ordinarily in the oil and chemical industry, the total amount of labor that goes into building the equipment is of the same order of magnitude as the total amount of labor for running the equipment for twenty years. If the price of the product falls

below the operating cost of the installation, that installation, and consequently that half of the total labor that built the installation, cannot be used for anything else. One reflection of this is that a large proportion of all the middle-class people employed in the oil and chemical industries are accountants, specialized in estimating financial outcomes into the far future under various long-term uncertainties about costs of raw materials, labor, and product prices.

Ordinarily the response to random variations in the raw materials, in the degree to which a particular wall is straight up and down, in what the latest client to walk in the door seems to have on his or her mind, in just what about quadratic equations a particular high school student seems to be having trouble with, is organized into an entirely oral information system—the informal world of skilled manual work, sales work, or classroom teaching. So the shortest temporal abstraction level is reflected in an information system within the individual, or at most in the interaction of a small group, and appears in our organizational analysis as human capital, as "skill."

As one goes up the hierarchy of most organizations, the time span of responsibility increases (Jaques [1956] 1972), which means in turn that the uncertainties one is responsible for responding to develop over longer periods. While the skilled worker or salesman responds from minute to minute to random variation, the manufacturing division responds to shifts in demand for automobiles that are updated each week, and the general office responds to quarterly or yearly profits, yearly shares of the market, and other indicators of where long-term investments should be made or where shifts should be made among executives who are not, evidently, responding effectively enough to weekly variations in demand. Research and development departments tend to be oriented to the very long run, though they are occasionally brought back to earth by a computation of profits on recent innovations by the central administration (Lawrence and Lorsch 1967; for critiques of this research, see Tosi, Aldag, and Storey 1973; Blau and Meyer [1956] 1987, 111–116).

The result is that information-processing systems have to aggregate information over different time spans, and when different parts of the organization respond over different time spans their information systems need to be segregated. We discuss in Chapter 5 how information systems for feeding market information back to engineering and manufacturing have to be faster and richer for innovations than for products from which the bugs have already been worked out. Consequently, the information systems for an innovation in general have to be segregated into an autonomous "divisional"–type structure.

Similarly, in Chapter 3 we discuss how the "well report" for past oil wells and the "well plan" for a well to be drilled in the same field bear an intimate relation over the very short run from the point of view of drilling operations. But all that absolutely critical detail disappears from long-term plans for connecting drilled production wells to the processing apparatus that prepares the oil and gas for shipment, for here the crucial question is whether the connections between the two have any leaks that could create a danger of explosion. From the point of view of commissioning the processing operation (starting production), the dates and mechanical details of the "completion" of the well are all the temporal detail one needs about drilling.

The level of temporal abstraction of any information system, then, serves as empirical support for an argument that the system is supposed to deal with an uncertainty developing at that time scale. Conversely, again we can assume the functional connection and use the level of temporal abstraction to find what sorts of constraints a given piece of the administrative apparatus is oriented toward—the constraints that operate at the time scale of the abstraction of the system. Such information flows with different levels of temporal abstraction are central to defining "decentralization" in Chandler's work, analyzed in Chapter 4.

4a. *Trustworthiness and Error Estimation.* Scientific papers are required to present the evidence on which the argument they make is based. Textbooks do not generally have to analyze the evidence, nor do they have to say what sorts of errors such evidence is subject to and what protections against error one has built into the experimental procedure. This difference is due to the difference in the degree of uncertainty about the arguments in question. In general, no competent chemist would disagree with any argument in a competent textbook in chemistry (for an exception, see Watson [1968] 1981, 110, 122). Scientists do not have to give their readers the basis on which to disagree with their reading of the evidence unless there is some uncertainty about how to read the evidence. It is assumed that if something is sufficiently uncertain to be worth publishing a paper about in a scientific journal, then it is sufficiently uncertain that one has to specify the evidence on the basis of which one makes a particular argument about it. Consequently, there are routinely descriptions of experimental procedure (in highly codified form, of course) in the scientific literature but not in the typical textbook.

A system of discourse that attaches qualifying information to assertions, so that the reader can go behind the assertions if necessary, according to the degree of uncertainty generally thought to attach to the conclusions, is a "trustworthy" system of information. Science, then, is a finely tuned trustworthy system because the relatively certain assertions

of the textbook have little data that allow the reader to assess the uncertainty, exactly because there is not much uncertainty; the more problematic assertion of a new discovery has to have evidence presented and evaluated in the standard scientific way. Hence, the detail of the evidence is proportional to the uncertainty of the assertion. Such systems do not provide useless information about certainties, nor do they present uncertainties as revealed truth.

This does not mean that there are no errors in the information. If there was no chance that scientific papers were wrong, scientists could just publish the conclusions without the evidence. And of course, the textbooks of a generation ago have in them things that were consensual in the scientific community at the time but that are now "known" by consensus to have been disproved by research in the meantime. A trustworthy system of information, in short, is one that does not routinely deceive its users about the uncertainty of the conclusions reached on the basis of that information but at the same time does not burden them with information about stable features of the environment. One does not ordinarily watch the news from the scientific community to see whether a given paragraph in one's physical chemistry textbook is wrong; one does watch that news to see whether conclusions of experiments reported in the most recent scientific literature stand up to different procedures in different laboratories.

One is not surprised to have a Nobel Prize winner write a sentence like "The letter was not in the post for more than an hour before I knew that my claim was nonsense" or for him to be advised by a mentor, who had forwarded some problematic work for publication, that "in this way, I [Watson] would still be young when I committed the folly of publishing a silly idea" (Watson [1968] 1981, 110). Scientific research is not worth doing unless it is in an area where the right answer is uncertain, and consequently scientific work is required to give others the basis for assessing its own degree of uncertainty. It is exactly because Watson was working in an area where he was quite likely to be wrong that he was required to give the evidence and the details of the argument.

Similarly in the oil business, everyone knows that many "promising structures" on the basis of acoustic exploration have no oil or gas in them and that one knows for sure only when oil and gas comes into an exploratory well. Further, one has a reasonable estimate of the recoverable reserves only after several exploratory wells have "delimited" the reservoir. A report of acoustic exploration thus has attached to it a professionally well understood degree of uncertainty, the first exploratory well a different (smaller) degree of uncertainty, and an estimate of recoverable

reserves based on several exploratory wells a still smaller degree. No one professionally concerned is likely to mistake one report for the other. The information system is "trustworthy," even if uncertain to different degrees after each estimate, because each estimate has attached a professionally understood assessment of its uncertainty.

If, however, an incentive system exists in which the executive in charge of estimating recoverable reserves is punished severely for producing an overestimate, while no serious consequences flow for him or her from an underestimate (presumably because one does not know one has lost money by not investing but always knows one has lost money by investing), that executive may not be rewarded for attaching the correct degree of uncertainty to the estimate. Then he or she may report lower reserves than the evidence justifies. The system becomes untrustworthy at that point, because there is a systematic bias in the estimate of the probable error. (It may be deliberately arranged to be biased downward if it is known that the decision based on the information will be taken by amateurs and so will not take account of the evidence on its degree of uncertainty; it would, however, be an unusual—and probably short-lived—oil company in which this was the case.)

4b. *Agency Problems in Information Trustworthiness.* One can tell that the travel voucher and entertainment expenses accounting system is oriented to contingencies of fraud and extravagance by the safeguards built into it. It is very hard to collect information on entertainment expenses that tells accurately how much entertainment expense it was rational to expend on the company's behalf, that provides a reasonable estimate of the marginal productivity of the last dollar of entertainment expense. In this case the low trustworthiness of the information is due to the fact that the person who has to report on the value of the entertainment to the company is also a person who benefits from a more expensive evening out.

Organizations write multiple rules about the external features of entertainment and travel expense (those governing travel by federal employees are a wonder to behold) that, one hopes, limit the degree of extravagance and fraud. Protections against error that an organization builds into its information system tell a lot about what the organization really wants to know.

Many of the sources of untrustworthiness of information are faults in the incentive system for providing accurate information and accurate estimates of how likely one is to be wrong. These are generally dealt with theoretically under the heading of "agency theory," which deals with the broad class of problems in which the agent has more information—later

or better news—than the principal about the uncertainties involved in a decision that the principal has to take (or that the agent has to take on the principal's behalf).

Agency theory deals with such things as "moral hazard" in insurance, in which the policyholder has information about strategies for reducing loss that the insurance company does not, but is not motivated to use those strategies because the insurance company pays the losses; or "adverse selection," in which the buyers of insurance policies know that they are worse risks than the company has estimated, so the insurance is a very good bargain for them, while those who fail to buy policies know that they are very good risks and so can better self-insure.

In general, information sources are people, who have motives to distort information so that they look better or worse than they really are. Such motives make information untrustworthy. The multiple devices that insurance companies have developed to deal with these problems (see Heimer 1985a, 11–17 and passim) indicate that trustworthy information about expected losses requires special structures to shift the incentive system. These devices also indicate that getting accurate information on the risks being insured, against the possibility of moral hazard or adverse selection, is a crucial uncertainty for insurance companies.

As Shapiro (1987) has pointed out, devices to make information more trustworthy can be simulated to make fraud more effective (for case material and evidence on how hard it is to find fraud, see Shapiro 1984). For example, a scientist's conclusions look more certain than they are if he or she reports faked data than if he or she merely reports an "expert opinion." Simulating the devices that we use to make scientific reports trustworthy makes false results look more trustworthy. Every device for making information trustworthy provides an additional reassuring cover for a con game when the device can be simulated; financial insurance companies with slim reserves may have the initials FDIC, for example.

A baroque structure of reviews, signatures required, committee consideration, and auditing usually indicates that the information system is endangered by agency problems, problems of motivated distortion of the content of an information flow; the tenure process in a leading university is a good model for what an information-processing and decision-making system with severe problems of agency looks like. But those same structures, by the functional argument, are indicators that agency-induced errors in the information will tend to lead to important errors in the decisions.

5. *Expertise of Analyzers.* The maintenance of the quality of the flow of information is in general achieved by constant minding by people who

know what uncertainty is to be analyzed, where the news about it is to be got, what causal unities determine which units are subject to that uncertainty, how to make the tradeoff between currency of information and noise reduction, what temporal span the decisions cover, what degree of uncertainty different sorts of information are subject to and how to indicate the level of uncertainty correctly, who is motivated to distort the system, what auditing or control procedures will work, and who are the other experts to consult on all these questions. Data on the expertise of those who mind the information flow supports a functional argument about what sort of uncertainty, and which constraints on success, the information flow was developed to control. Conversely, assuming the functional argument, one can tell that if high-priced accountants are minding the information flow, the long-term financial outcomes are the crucial things being estimated; if engineers in charge of maintenance are minding it, the constraints are likely to be reliable performance of machines; and if auditors whose chief qualification is suspiciousness go over the accounts frequently, the validity of claims and the likelihood of fraud is a crucial constraint on success.

Structure and Function

Many functions never get fulfilled: getting accurate information on the quality of a used car when the place where the news is located is the used car salesman is not a social structural possibility. Conversely, many functions that are carefully fulfilled solve problems that were never there in the first place: certifying that high school teachers understand curriculum theory or audiovisual methods of instruction has a clear manifest function, though it seems to do no good (the movie projector in the school gives both teachers who have and those who do not have audiovisual training a few minutes peace and classroom order, a true and important function).

Functional arguments are always suspicious because they attempt to explain things by their consequences, and this obviously requires a reverse, consequence-governed, chain of causation, a causal loop so that a consequence can explain its cause. Such arguments are particularly suspicious when one is explaining the structure of information flows in a situation in which rewards depend on what information is given. Teachers and teacher professional organizations, through their certificates, create an information system that looks as if expert judgment is used to select instructional movies. Of course, the one way to select good instructional movies is actually to look at a lot of movies in the subjects

one teaches and pick out the best ones. But no school board (that I know of) is willing to pay teachers a skilled worker's wage to watch movies all summer.

The function of the certificate is to monopolize positions as teachers for those who have gone to education schools (where they teach audio-visual methods) and completed liberal arts programs, and to provide something to justify having people around that education school who produce or market instructional films. Since the manifest function of selecting good instructional movies would be very expensive and no one believes in such selection enough to pay that much for it, and since the latent function of this certificate is important enough to some people to keep the system going, an information system is created with very little information in it (certification systems for skill are treated as information systems in Chapter 6; the origins of skill requirements are treated in Chapter 2).

In this case we can perhaps identify the causal links back from the latent function to the structure and content of the information system, having to do with the fluctuating power of education schools (which have more luck requiring such certificates when teacher training is popular and many teachers are required by the system). Consequently, we do not have to depend on the manifest function for an explanation.

But the first evidence we used to suggest that all was not as it seemed was to check whether the information system had the characteristics that would be required for its manifest function: is the information in the system about the principal source of uncertainty on instructional films—namely, whether the film is appropriate to the students and to the instructional purpose of a particular teacher? Since the judgment of whether a film fits is clearly best made by a teacher seeing the film, the characteristics of the required information system are clear. One can very often disprove functional hypotheses without examining the reverse causal chain, by examining internal evidence about what the structure would have to do for the functional hypothesis to be true.

The field we are mining in this book, however, is shot through with situations in which people's rewards depend on the information flowing in an organizational system. Because a sociologist's fundamental inclination is to look at the seamy side of life for explanations, the big sociological fact may seem to be that people are motivated to misinform. To convince the sociological community that explanation by manifest functions of getting the organizational job done is fruitful is therefore rhetorically difficult. Our main hope is to do so by showing that many details about the setup of organizations that are otherwise inexplicable

are explained by the proposition that information systems tend to grow toward sources of news about the central uncertainties of that part of the organization, and that the structure of the information-processing system can often be predicted from the assumption that its design gets the job done. Thus, the alternative hypothesis that the information system is a sham, designed to give cushy jobs to its perpetrators, can sometimes be refuted by showing that that alternative does not explain the elegant fit between what an information system does and what is needed to get the job done.

A second problem with functional explanations is that they take such things as organizational constraints, and uncertainties bearing on those constraints, as real, as causal variables. These causes work not only by functional mechanisms, already suspicious, but at a supraindividual level of explanation as well. Since, of course, all the significant organizational acts that we will be explaining are done on purpose by conscious individuals, any explanation at the social level has ultimately to be shown to be adequate at the level of intentions of individuals. I address this problem in detail in Chapter 4, in connection with Alfred D. Chandler's (1962) analysis of the origin of the multidivisional structure in Du Pont. But we can sketch here the broad lines of attack that this book takes.

The first line of attack is that in fact, empirically, we often find that an organized structure of intentions of a lot of different organization members exists, and that that organized structure is such as to get the organizational job done. The interconnections of those intentions are of such a character as to cause organizational information systems to grow toward the sources of crucial uncertainties for organizational success, and the substantive content of the structure of the information processing system as created by those intentions is in fact such as to collect information on the right units—to reduce noise, to abstract in the right temporal frame and to segregate different temporal features of uncertainties, to assess the degree of uncertainty accurately and to control agency problems of information quality, and to put the right sorts of experts in charge. If in fact individuals make up functional structures, this is surely an important dependent phenomenon that people who specialize in explaining individual behavior ought to address. This individual-level explanation is God's work, but not my own particular vocation.

The second line of attack is to assume that people as individuals confront the problems posed to them in their individual work, and that a long series of solutions to such problems of how to do one's individual assignment are such as to make the organizations "climb the gradient" and grow a structure that responds to the crucial uncertainties that

govern organizational success. We assume organizations climb that gradient faster if the people in charge of creating and maintaining information are truly experts, if the general organizational tradition is one that discourages corruption, fraud, and manufacturing evidence and the like. This leaves the intellectual challenge of specifying how organizational problems become problems assigned to individuals who have the "obligations, rights, incentives, and resources" (Heimer 1986b) to solve them.

A third line of attack is an evolutionary one, particularly relevant to Chapter 5 where we specify some features that administration of innovations needs to have. A careful reading of that chapter should suggest that many organizations are not well set up to introduce innovations, and may fail even if they are because introducing innovations is organizationally difficult. The uncertainties that confront an innovation are severe in all sorts of areas. But the high rate of failure of innovations means that the features we find in those organizations that innovate successfully are likely to be functional for innovation.

If in the long run business firms can be profitable above the interest rate only if they introduce innovations (as Schumpeter 1942 argued), and if organizations last longer if they are profitable, then organizations whose innovating arms have grown faster and more accurate information systems that reduce the uncertainties connected to the innovations are more likely to be around for sociologists to study (cf. Nelson and Winter 1982). Similar evolutionary arguments may be used to suggest that administrative structures in General Motors, Du Pont, or Sears may be more functional than those of Kaiser-Fraser, Lydia Pinkham, or Montgomery Ward.

The Plan of the Book

The overall organization of the book is to ascend the levels at which information and decision systems are organized, starting with individuals' skills and ending up with class relations in whole societies. We pass from individual skills to technical departments in manufacturing, then to divisional organization, to innovations that cause reorganization of these first three, to networks of contracts among organizations, to segmentation of labor markets in the economy as a whole, and to the formation of workers' information on their class and political interests. We then hop sideways from these materials (largely focused on economic organizations and their offshoots) to illustrate the application of the theory to university research administration.

In Chapter 2 we start with variations in the basic information-pro-

cessing mechanism of modern social organizations: the employee. The main social structure to be explained in this chapter is the division of labor, the main information-processing mechanism is skill of individuals, and the main social effect is the stratification of employees by level of skill. This definition of the problem is implicit in Charles Sable's *Work and Politics: The Division of Labor in Industry* (1982), so we base our analysis on that book. Our main problem is to specify precisely enough what skill consists of so that we can build a theory of when it will be necessary.

We argue that all work is characterized by an inherent dilemma: that productivity is highest when all the activities necessary to production or delivery of a service are highly routinized but specialized human work is most valuable (and hence the incomes of organizations highest) when there is uncertainty about what to do. Routinized work is more productive partly because people work faster when they are doing something that is routine for them, but mainly it is because, once routinized, the work can be mechanized or computerized or otherwise improved, and the improvement will stay in the productive system as part of the routine rather than as a sporadic flash of genius.

But routines are in general effective only if used under the appropriate circumstances: only if the parts to be assembled are standardized, only if the clients in a restaurant come in with standardized wants and select from a limited menu, only if the pilot can depend on his (or her, if airline piloting were different) counterpart workers in the control tower to behave in a routinized fashion so that the information they give him has an absolutely clear meaning. This means in turn that uncertainty in the environment of work—unstandardized parts, unstandardized clients, or unstandardized fellow workers—undermines productivity by undermining routinization.

We will analyze skill as *the capacity to routinize most of the activity that comes to a given work role in an uncertain environment.* We will argue that skill is a repertoire of routines which the workers can do accurately and fast, as well as a set of selection principles among routines, such that the complex of routines and selections among them deals with most things that uncertainty brings to the worker. Thus, we will expect to find skill when a great many different things must be done to produce the product or service but when each of those things has to be done in several different ways depending on the situation. For example, more complex products are likely to require higher skill levels in the work force that produces them, but the highest skill levels should be found in industries with complex products that have to be customized for particular clients and constructed in varying environments, as in the building construction

industry. Much of the chapter, then, consists in the analysis of variability and uncertainty in exactly what work has to be done at a particular time, uncertainty about which of the routines in workers' repertoires have to be pulled out.

Chapter 3 turns to the variation in structures of administration in different sorts of activities within manufacturing, by showing that the system for extracting, treating, and shipping crude oil has many different types of information problems, and so many different structures for processing the information. The data for this chapter come from a study of building such a manufacturing administration in the Norwegian State Oil Company (Statoil), the first such organization built by a Norwegian oil firm (other than the Norwegian branches of multinational oil firms, which of course are legally Norwegian oil firms). The different things one has to do, say, to build an information system adequate to buy a stock of spare parts rather than to develop accounting software or to drill a production well a couple of miles deep without killing anybody are especially clear. (Much of the material for Chapter 3 has been previously published in photocopy form by the Institute of Industrial Economics in Bergen, Norway [Stinchcombe, 1986d], which sponsored the research.)

Chapter 4 reanalyzes Alfred D. Chandler's classic *Strategy and Structure: Chapters in the History of the American Industrial Enterprise* (1962). The central concern of that book, we argue, is how information coming from different sources, especially different markets, is combined to make different sorts of decisions. Chandler argues that when the information needed to make money off a company's products or services comes from several sources and has to be related to several distinct systems of decisions, one has to form more or less autonomous divisions to integrate the information from a given market and then arrange for the division so created to pass on to the central office only abstract and general information relevant to financial and investment decisions. Chandler's argument then becomes a functional one: that *in fact* organizations that need divisional decentralized structures because they are in multiple markets tend to get them.

First we undertake to define more exactly, in terms of information-processing and decision-making structures, what Chandler means by his dependent variable, decentralization into divisions. We then examine in detail how Chandler builds this functional argument for a case where he has the most historical detail, the origin of the divisional structure of Du Pont. Chandler has a good deal of information on how individual actions inside Du Pont added up over the course of several years to a shift to a divisional structure. This gives a hint of how the methodological

objection to the organization-level functional analysis (that is the basic strategy of this book) can be disaggregated into individual human action that is understandable.

The core of the analysis of Chapter 4, however, is the explication of Chandler's implicit definition of when a firm is in multiple markets. Our argument is that, for Chandler, a market is a matter of where the information about the central uncertainties about a flow of products or services lies, and how this information can be got to decision points and integrated in a sensible way, rather than primarily about who the firm's competitors are. Chandler uses the economic competition definition of what a market is only occasionally and tangentially; his main, implicit definition is that a market is a number of phenomena outside the organization that produce a common administrative problem—problem of relating information to decisions—within the firm.

In one sense, then, the subject of Chapter 4 is the origin of the overall variation among firms between divisional and centralized organizations. But the core explanatory principle is that *within* the firm the problems of different product lines are different enough to need different information-processing structures. One needs to separate into different divisions administrative dealings that involve different sorts of uncertainty, because different product lines have to attend to different sorts of news, collected in different ways, evaluated according to the particular situation of that line of goods, and integrated with engineering and manufacturing in a distinct way. One needs one information-processing system and set of routines for making decisions to sell explosives to a few mines or the War Department, and another for making decisions to sell paint to thousands of householders or housepainting firms. The central administrative mistake, according to Chandler's argument, is to administer news about one sort of uncertainty with a structure built to handle another kind. Thus, Chandler's argument is exactly in tune with the main thrust of this book.

Chapter 5 deals with the special kinds of uncertainties that are ordinarily associated with innovations and with the sorts of information-processing and decision-making problems that they pose. The analysis can be seen as an elaboration of Joseph A. Schumpeter's examination of the relation between economic innovation and routine administration (e.g., Schumpeter 1942). Schumpeter identified routinized administration with large bureaucratic organizations. But in order to have the economic effects Schumpeter urged, innovations must be produced by efficient—which means routinized—production processes and must reach the market in routinized channels. The temporary monopoly over

a valuable product or service that in Schumpeter's analysis creates the profits from innovation does not in fact produce profits unless the innovation can be produced on a large scale and marketed to collect the extra margin created by the monopoly.

Our argument will be that the introduction of an innovation involves a higher level of uncertainty than do the production and marketing of goods that have been on the market for a long time. The news needed to improve the innovation, to find its market, to introduce elaborating innovations, and to penetrate new niches is likely to be found in a different place than those which the information system of an existing organization reaches. And it has to be dealt with simultaneously and rapidly by engineering, manufacturing, and marketing in much the way that Chandler argues for distinct product lines.

Thus, Schumpeter is right in stating that it should ordinarily be hard to introduce innovations by using an existing bureaucratic structure. Yet a structure with many of the characteristics of a bureaucracy is needed to get the profits out of the monopoly created by innovation. An innovation poses an administration with a problem similar in form to the dilemma that we argue in Chapter 2 produces a demand for skill. It also poses a problem of differentiated information collecting and decision making similar to that which Chandler argues produces divisional decentralized administrative structures, as we argue in Chapter 4.

The sociological answer to Schumpeter's argument about the unfitness of bureaucracy for innovation (aside from the empirical observation that most important innovations are produced and marketed by large bureaucracies) and the consequent necessity for an individual heroic entrepreneur is that the administrative problems posed by innovation are of the same sort that skill and decentralization ordinarily solve, though perhaps more extreme. Thus, the first functional argument here is that innovations should tend to create pressures for divisional decentralized administration, such that the innovation has a separate but integrated news-collecting, information-processing, and decision-making structure. The second functional argument is that, because the level of uncertainty of the work involved to introduce an innovation will ordinarily be higher than that of other production, such a division is likely to have a higher skill mix than a division with an equally complex product or service that is no longer an innovation.

A good portion of Chapter 6 has been published previously as "Contracts as Hierarchical Documents" (Stinchcombe 1985b). The argument is that many contracts between organizations are actually formal organizations themselves, with news collection, information processing, and

decision making built into a joint social structure consisting of the two (or more) corporate parties to the contract. Thus, the contract creates a social structure to deal with uncertainty, just as the labor contract does. The chapter then challenges Williamson's argument (1975, 1979) that dealing with uncertainties produces hierarchy. (Williamson's later work [1985] is compatible with the argument presented here.) The argument is that anything one can do with social structures constructed by labor contracts, one can do with social structures constructed by contracts between firms.

Williamson's argument is of the same functional form as ours, namely, that when information and decision problems create certain kinds of uncertainties, then a social structure that tends to resolve those uncertainties (that is, a hierarchy) tends to grow to manage them. The only problem with his argument, we will allege, is that the functional requirements of hierarchy itself involve continuity of production, such that the heavy investments in building a hierarchy and training the people to do their roles in it will be productive long enough to pay for themselves. Consequently, when one is only going to be, for example, in the business of building oneself a factory or an office for a couple or three years, one does not want to erect a hierarchy able to build factories and offices on into the twenty-first century. Rather, one wants to hire organizations that know how to build factories and offices, and that have skill structures and incentive systems appropriate to the uncertainties of that business, to put together a temporary social structure to build the building. Then, if it turns out that the site is sand rather than rock, one wants to change some of the specifications for the building in the middle and still use the same system of contracts and subcontracts to build the new design.

We will show that when, for one reason or another, it is hard to build a hierarchy to do the job but when the problems that (Williamson argues) tend to produce hierarchies occur, the structures of hierarchy will tend to be built into the contract itself. Thus, we are supporting Williamson's functional argument by breaking down the variable of "hierarchy" into its components, so as to avoid identifying it with the legal unity of a firm or another organization. In short, when functional pressures toward hierarchy exist but there is difficulty building an ordinary hierarchy made up of employment relations, we find instead contractual means for creating hierarchies among corporate actors.

Chapter 7 returns to the problem of skill of Chapter 2 from a different point of view. If the division of labor in a set of activities produces a set of skills to deal with the uncertainties, as outlined in Chapter 2, then the personnel system of the organization has to create a status system

to motivate the acquisition and utilization of those skills, as well as a recruitment system to fill the jobs with people who will turn out to have the appropriate skills.

One of the most complex uncertainties that organizations deal with is uncertainty about whether people will be willing and able to do the work. The information system inside the individual is just as complex as the information system of the organization. Measuring what that system is capable of, then, involves analyzing the relationship between the particular complexities and uncertainties of the work to be done and the complexities and skills of the mind and body of the worker or recruit.

It turns out to be very difficult even to measure how well people are doing their jobs because, in general, line management does not fully know the complexity of the workers' jobs, and personnel management knows even less. Unless one knows a job very well, it is very hard to differentiate explanations of low productivity that blame true variations in the world from those that involve lack of competence or will on the part of the worker. But even when one can measure performance accurately because one knows the uncertainties involved in the work well, it is very hard to predict what quality of skill and what discretion in the internal decision-making system a worker will develop over the course of his or her years of experience with the organization and, hence, what kind of productivity, in the environment to be confronted over those years, the organization will get. Further, it is exactly this sort of information about employees' performance that is hard to collect from the people who know most about it—namely, the workers themselves—because they have an interest in getting the rewards that come from giving good news about the uncertainty about their competence and avoiding the punishments that come from giving bad.

We deal with four main structures for dealing with the fundamental information uncertainty of the labor market: continuity in the job (seniority in a job as a basis for holding that same job), internal promotion systems ("internal labor markets" in the strong sense), certificates from schools, and certification by a body of peers. We argue that reliance on such imperfect information systems provides the structural basis for segmentation of the labor market, because the fundamental privilege in the labor market is the capacity to give reliable certification that one can do a job. We do not, however, show why it is that some groups (e.g., white males) are more able to give such certifications than others (e.g., blacks or females).

Chapter 8 uses the materials from Chapters 2–7 on how organizational structures originate to examine the problem of class consciousness.

It is very clear that working-class consciousness is a consequence of the formation, with industrialization, of new kinds of organizations in the economy. Much of the argument in Chapters 2–7 can be conceived as an analysis of how capitalist (and socialist) economic organizations differ from peasantries, agricultural villages under feudalism, and other preindustrial forms of organization. In particular, we argue that such reorganization of the economy tends to produce reorganization of the labor contract, such that the *same* labor contract is imposed on whole categories of people. Labor contracts in preindustrial social structures tended to be more individualized.

From this point of view, class consciousness is a product of collective contracts. At first, collective contracts are unilaterally imposed on the workers by urban industrial and craft employers, creating categories of workers subject to the same working conditions, the same measurement of performance, and the same wages. Worker class consciousness is simply organizing those categories to negotiate their side of the collective contracts already imposed on them. This is a caricature of the argument in David Lockwood's *The Blackcoated Worker: A Study in Class Consciousness* (1958); a more nuanced development is given in Chapter 8.

But class consciousness, in the way it is ordinarily used, is the projection of one's position in an employing organization onto the society at large, so that one's place in that society, one's interests in the economy and polity, are seen "as through a class darkly," in Kenneth Burke's phrase. We therefore reformulate E. P. Thompson's great book, *The Making of the English Working Class* (1963), to specify under what conditions a position in an employing organization will be projected onto the larger canvas of labor market organization and political movements as a definition of interests there.

We then apply this combined Lockwood-and-Thompson theory to the decline of class hostility in the more class-conscious (industrial and transportation blue-collar) part of the economy and the low level of class consciousness in many parts of the service sector of that economy. Chapter 8, then can be thought of as an extension of an analysis in Chapters 2–7 that is oriented dominantly toward economic employing organizations, to the explanation of some features of unions in the labor market and of left political organizations. It does so by showing how the structure of organizations may provide information to workers about their interests, information that is then interpreted by the culture of the larger system.

Chapter 9 carries out another sort of extension, into a field involving employing organizations whose purposes are education and research. We

give a couple of case studies of administrative structures in universities, namely those that administer space and those that set teaching loads of professors.

A prestige university confronts the uncertainty about where new knowledge is to be found and who will find it. Clearly, unless a piece of knowledge has been highly uncertain beforehand, it is not much of a research accomplishment to discover it. Thus, as in professional athletics, research tends to be completely uninteresting if the outcome is certain and most interesting in cases where nobody knows for sure how a given line of investigation will come out. From an administrative point of view, the main implication of this situation is that the people who know both what work is to be done and who can do it are concentrated at the lowest hierarchical level of the organization—the professors—though they may be paid more than people with many subordinates. What these subordinates hopefully know how to do is to have good bets about where new knowledge is to be found, better than those of fellow scientists at the lowest levels of *their* organizations who provide the baseline of uncertainty.

Scientific results, and to a certain extent results in other scholarly disciplines, change the degree of uncertainty generally held in the scientific community. Those findings or theories that are most fruitful, in the sense that they most change other scientists' bets about where new knowledge is to be found, are the most valuable. That is why citations turn out to be such good measures of the subjective estimates of scientists of who has done the most important work.

But this deep uncertainty—in which other competent people cannot bet as well as the future Nobel Laureates one wants to hire, retain, and promote—makes management in the ordinary sense extremely difficult. Even the knowledge of what sort of work the organization will be doing in chemistry or physical anthropology is concentrated toward the bottom of the organization, and people in chemistry and physical anthropology at other universities can probably usually bet better than the dean what will occupy the space in the labs and the research time of the dean's subordinates next year and the year after.

But this means in turn that the information the dean needs in order to make organizational tradeoffs between departments, for example to decide on allocations of laboratory space or teaching loads of various faculty members, have to be wrested from subordinates who have no interest in providing the basis on which space might be wrested from them or in having students and courses added to their teaching load. Thus, the commitment of a university to research distinction is in essence a com-

mitment by the university administration not to know how to manage the organization, not to know whom to hire, retain, and promote, and not to know how to organize the flows of the most crucial information; their subordinates, the professors, and professors in other universities, are the ones who know.

In the administration of the allocation of space we analyze why this inherent situation of university authorities tends to lead to departmental sovereignty over space, with only occasional invasions from above. In the administration of teaching loads we study the pressure for university administrators to ignore measurements of individual research productivity in setting teaching loads, even though the dominant reason for having few courses in the (more or less uniform) teaching loads of universities is to give time for research. The question to be answered, then, is why universities so seldom try to allocate more time for research to those who do it best and why instead they depend on the National Science Foundation or the National Institutes of Health to make the judgments about whose teaching responsibilities should be lower, whose higher, by judging whose time those institutions are willing to buy from the university (at cost plus 50 percent or so—the percentage of overhead is higher at some leading universities).

The objective of this brief exploration into the organizational aspects of the sociology of science to show that uncertainty and the information needed to reduce uncertainty shape the structure of administration of universities as much as they do that of economic organizations.

The objective of the book as a whole is to illustrate in a variety of contexts how the social structure of organizations can be explained by the structure of the information problem they are confronted with. The idea is that organizational principles differ radically from one situation to another because what is functional for an organization differs similarly. What individuals have to learn; what distinct types of information structures must be built into manufacturing management; which departments processing which information should group together into divisions; how the administration needed for managing innovations and that needed for managing more routine products differ; how management tells whom to hire, retain, and promote; what the content of contracts among organizations will be; what features of organizational position will be taken by an organization's members to define their fundamental economic and political interests; and why scientists go their own way regardless of their administrators—all are to be explained by the nature of organizational uncertainty and the form of information processing needed to reduce it.

Broadly speaking, we are trying to explain why the formal structure of

organizations varies depending on what those organizations have to do. All the elements of this theory exist in the organizational literature. For example, the centrality of routines in organizations, emphasized in Chapter 2, is central to both Cyert and March (1963) and Nelson and Winter (1982). The dependence of skill levels on nonroutineness and the location of skills in those parts of manufacturing with much market uncertainty are analyzed by Charles F. Sabel (1982). But those pieces have not been put together into a systematic theory of the skill levels of organizations.

In particular, it has not been done in such a way that the same theory can be used to explain why innovations not only create such administrative difficulty but also require such high skill levels, as we do in Chapter 5. To make the transition, we have to use the complex theory implicit in Chandler's interpretations of various concrete business facts in *Strategy and Structure* (1962). But that theory is not systematic enough to use for our purpose. For example, Chandler writes as if the explanation for divisionalization of firms is the same as that for general office structure. But divisions respond to different uncertainties than does a general office, and so are differently explained. This distinction helps to untangle what Chandler's theory is and so to show why his analysis of Sears is not convincing.

Because Schumpeter had not read Chandler and had not studied Du Pont, he could not imagine making bureaucracies flexible enough to be innovators. He therefore placed all the burden of economic advance on heroic entrepreneurs. But Chandler's theory shows more exactly what kinds of bureaucracies can innovate, can turn an invention into a going concern making monopoly profits.

Thus, the work of this book is to make an overall theory of uncertainty and information sufficiently strong and exact to show where our geniuses in organization science went wrong. It does so by a systematic attempt to explain variations in the formal structure of organizations. These variations are often between parts of the organization. For example, to explain the structure of Du Pont, with product line divisions and a general office, we need to explain why decentralization to divisions involved first centralizing those divisions. But we also need to explain what uncertainties the general office responds to and therefore why it needed more abstract communication flows to it and a nonhierarchical committee structure of communication within it. The structure of the general office in Du Pont was a committee structure much like one that had been tried and rejected for internal administration of the divisions.

In some sense, then, the theory of this book is old hat. But by explain-

ing variations in organizational structure by variations in the type of uncertainty dealt with, we can show the relations among the classics of organization theory. In the process, we make the classical theory empirically richer and make it stretch from the micro level inside individuals as they develop skills to the macro level of class relations in whole societies, from the General Motors assembly line to the university laboratory for research in biophysics. We hope, in sum, to make the theory we have used both more concrete and more general, so as to explain why the forms of organizations are at once so wonderfully various and yet so obvious once we see the uncertainties that they have to deal with.