Chapter One

Language Matters

Sociological theorists rarely give explicit attention to the language used in formulating arguments and drawing conclusions. We argue that attention ought to be paid to the choice of a theoretical language. This chapter discusses such language matters; and it also argues that choice of a language matters for what a theory can express and whether (and how) it can unify fragmentary knowledge.

1.1 LANGUAGES FOR THEORY BUILDING

Sociological theories can be built and arguments can be unified without any special tools or methodologies. However, we believe that attempts in this direction face daunting challenges. The essential problem lies with language, with the natural language used in most sociological theory.

Natural languages such as English and Hungarian are less-than-perfect tools for eliminating ambiguity in the claims that one might wish to make. In fact, ambiguity normally comes in two varieties.

Lexical ambiguity. Although technical terms might be defined formally, the choice of words used in expressing their meanings brings about associations that soon generate deviations from the initial definitions. Such deviations lead, in turn, to ambiguities. We describe such ambiguity in detail in analyzing the concept of niche in Chapter 8.

Structural ambiguity. Even when a language is freed from lexical ambiguities, and it "sounds" like the language of first-order logic, there is still room for structural ambiguity. For example, the sentence "There is a period of endowment for all organizations" can be interpreted in two ways. According to one reading, the endowment period is the same for all organizations; a second, more liberal, reading allows for different periods of endowments for the different organizations. In Chapter 7, we make clear that the choice between these readings matters for the explanatory power of the theory, so that it should be settled prior to implementing a formal approach.

As a result of such ambiguities, theoretical arguments in natural language admit multiple interpretations. Historians of classical sociological theory have written hundreds of papers on what Weber, Durkheim, or Marx really meant in certain passages. Such very close study of the passages does not appear to yield agreement on the interpretations of the sociological classics. The situation does not appear to dif-

fer for contemporary theories, which suggests that the natural-language statements are not getting any clearer.

Such unavoidable ambiguity makes it very hard to analyze complicated arguments in a natural language. Questions inevitably arise about whether some combination of premises does or does not warrant a given claim (or yield some inconsistency). It is notoriously hard to reason about the implications of complex structures stated in natural language. The chances of success in building coherent theories run higher if the arguments are stated in some kind of formal language. This is the strategy that we follow.

Success in efforts at formalization depends critically on issues of interpretation and translation. Translation from natural languages to constructed (artificial) languages is far from automatic. Introductory courses and texts in logic seek to train students in making consistent translations of indicative sentences, sentences that state facts (e.g., "Durham University is an organization"). Once students have learned the rules of this game, they are confronted with more complicated indicative sentences, e.g., "Every organization faces a competitor," whose translation requires learning the mysteries of quantification. But even the initiation into these mysteries leaves some problems unsolved. Translating indicative natural-language sentences to a logical language has to clear the hurdles posed by lexical and structural ambiguities. In other words, the sentences have to be interpreted. Formal languages are designed to rule out lexical and structural ambiguities; but the evolution of natural languages did not follow such guidelines.

Sometimes the logical structure of claims in a natural language appears on the surface; but often it is hidden. "Every person has a mother" indicates a quantificational structure on the surface. But "Men are mortal" does not. Both sentences have a conditional nature; but the conditionality is not always given on the linguistic surface by typical expressions such as "if . . . , then . . . " Although a course in (first-order) logic might provide some guidance for identifying the logical structures of natural-language statements, one has to analyze the *use* of a sentence in scientific argumentation in order to excavate its logical structure. Finding the "intended" logical structure of a sentence requires interpretation. Sometimes it is sufficient to analyze the sentence to find such an interpretation. In other cases, attention to the theoretical context and examination of how a certain sentence is used in argumentation is needed to find the intended interpretation.

To this point, our examples share a tacit assumption of most introductory textbooks in logic, that the intended interpretation is available in logic (typically first-order logic). That is, the standard texts assume that the expressive power of the formal language is sufficient to distinguish among sentences with inferentially different behavior. But this need not be the case in real problems. For example, "Men are mortal" and "All men are mortal" might appear to be similar from one point of view. But in conjunction with "Socrates is a man," the second sentence leads to the conclusion that "Socrates is mortal," whereas the first sentence does not. "Men are mortal" might be sufficient to support the conclusion "Socrates is presumably mortal" in an appropriately chosen logic. Classical first-order logic does not have sufficient expressive power to account for this difference. Nonetheless there are

logics with sufficiently fine-grained semantics to deal with this challenge.[1]

Given an infinite variety of logics,[2] no formalization can avoid the fundamental problem of choosing a formal language with adequate expressive power. Some typical problems in this respect can be analyzed with reference to the notion of extensionality.

Extensionality means that the replacement of names by other names with identical denotations (or predicates by other predicates with the same extension) yields logically indistinguishable expressions, that is logically equivalent formulas. Extensionality simplifies the logic, but it makes its notion of meaning rougher. In other words, the notion of logical consequence in first-order logic is less sensitive to nuance than are intensional logics (logics that relax the restriction to extensionality). Classical first-order logic, for example, cannot model the inferential impact of modalities such as perception or default, nor can it express the quantificational structures of rules with (unknown) exceptions, such as "Men are mortal."

Our study of relevant sociological theories suggests that considerations of genericity and modality arise regularly. Linguists refer to bare plural sentences such as "Organizations seek to seal off their technical cores from environmental uncertainty" as generic sentences, because they are general but not universal. These sentences typically express generality with (possible) exceptions. Modalities are expressed by sentential operators such as "possibly," "necessarily," and "allegedly." These modalities (or attitudes) do not fit into an extensional framework. Because genericity and modality appear to be ubiquitous in sociological theorizing, attempted translations of these theories into standard mathematics (whose underlying logic is the classical first-order logic) might not capture some key intuitions. Indeed, this problem might explain why formalizations in sociology so often seem to sterilize arguments, washing away the insights that made them appealing in the first place (Hannan 1997a).

A central challenge in formalizing and integrating sociological theories is to retain the core insights in making the translation from natural to formal languages. This task can be especially hard if the formal language is some kind of classical mathematics, such as algebra, calculus, or probability theory, as just noted. Fortunately, modern logic offers some appealing alternatives. The syntax and semantics can remain reasonably close to the natural-language argument; issues of genericity and modality can be handled systematically; and one can still "calculate," that is, derive implications, prove soundness, check for consistency, and so forth. For these reasons, we base our theory-building strategy on the use of modern logic(s).

In addition to choosing a language, a formalizer must adopt a strategy. A common misconception about strategy in formalizing theories comes from the mode of presenting formal work. The usual published report of a formal analysis begins by laying out definitions and assumptions and then uses them to derive lemmas and theorems. This mode of presentation suggests a mechanical process whereby deep insights follow from innocuous assumptions. Although this picture perhaps reflects

[1] We sketch such a logic in Chapter 6.

[2] The use of "infinite" in this context is not hyperbole. Modal logic alone contains an infinity of defined logics.

a common belief about scientific activity, it seems hard to find cases of insightful theories that were actually generated in this way. In particular, insight seems to have no role in this view, and there is certainly no guarantee that anything derived in this manner will be insightful.

Because they more closely resemble scientific work as we know it, we prefer two alternative views of theory construction. We attribute both views to Imre Lakatos (1976, 1994). Both hold that insights occur first (in any imaginable manner) and then formalization is undertaken as a way to preserve, strengthen, and possibly amplify the impact of the insight.

Lakatos' first alternative sees the central insights of an argument as stemming from its conceptual definitions and the explanatory principles (causal stories) that link them. Here the formalization effort seeks to clarify how and when the explanatory principles operate by recasting them in a sharper language. As a result, the explanatory principles inherent in the insight might be discovered to be more or less consequential than originally thought; they also can be more readily compared to accounts based on other rival principles that attempt to explain the same phenomena. We illustrate this approach thoroughly in chapters on age dependence (Chapter 7) and the niche (Chapters 8 and 9).

The alternative view holds that the deep insights reside in the key theorems. Given a theorem containing such an insight, the formalization enterprise attempts to identify assumptions that might be used to derive it. In other words, the formalization effort works to create a formal argument that yields this "target" theorem as an implication. In the process, some assumptions might be found to be more plausible than others, just as some sets of assumptions might be more parsimonious than others. The chapters that analyze density dependence (Chapter 4), resource partitioning (Chapter 9), and structural inertia (Chapters 10 and 11) follow this strategy. Other chapters typically use a combination of the two approaches.

As these two views suggest, formalization does not automatically generate insight. Rather, it allows us to understand and sharpen insights already in mind by tracing ideas and concepts through chains of reasoning that might be hard (or even impossible) to follow precisely in natural language. For instance, one possible result of the target-theorem approach is that the formal analysis might show that a theorem does not follow from commonly accepted or empirically verified premises. On the other hand, if the target theorem can be shown to be implied by sensible premises, then the process of building the formal model might lead to unexpected insights as well. For example, fewer or different assumptions might be needed than was initially believed. Or the formalization might produce additional, previously unconsidered derivations that could be subjected to empirical scrutiny.

An important related point is that formal tools—including logics—can often be used to show inconsistencies or deficiencies in verbal theories that contain nontrivial chains of reasoning. This kind of demonstration makes little contribution, in our opinion. After all, a main reason for using logic and mathematics in science is to avoid the ambiguity of natural language. So no one should be surprised to learn that a "rational reconstruction" and analysis using mathematics and logic can uncover the deficiencies of natural language as an analytical tool. The greater challenge in using logic and other formal languages—the one that top modelers take on—is to

deploy the formal tools to deepen the insights contained in a verbal theory, even if this takes considerable reconstruction. As the theory building in this book illustrates, reconstructing a verbal theory in a formal language is rarely a trivial, or even an easy, task.

1.2 USING DYNAMIC LOGIC

In the sociological context, an unusual feature of our work is its explicit use of modern logics. We now explain why we make this strategic choice.

The main role of logic in scientific applications is to clarify the notion of "logical inference" or that of a "sound argument." Many logics have been proposed over the last 2500 years. Most of them share a core idea: an argument is sound if and only if it is impossible that all of the premises are true and the conclusion is false. In other words, an argument is sound if the truth of the premises guarantees the truth of the conclusion. Given that various well-studied logics share this view, how can they differ? These differences arise from many sources. For instance, some logics construct and use different formal languages, others use somewhat different conceptions of truth, or different interpretation of what "impossible" means.

The first (more-or-less) formal system of logic, Aristotle's syllogistic logic, focused on reasoning about premises that express some kind of quantification. That is, it analyzed logical implication for sets of sentences that contain expressions such as "all," "some," and "none." (Recall the paradigmatic syllogism: "All men are mortal"; "Socrates is a man"; therefore, "Socrates is mortal.") Another Greek tradition of logic, that of the Stoics, dealt with sentential connectives such as "not," "and," and "if . . . then."

Today logics similar to the Aristotelian and Stoic logics are called propositional logics, because they do not deal with the internal structure of the propositions. Modern logics, beginning with predicate logic (also called first-order logic) deal with quantification and sentential connectives and also analyze the internal structure of propositions. In a sense, predicate logic incorporates (modern) propositional logic as well the (modern) logic of quantification. Modern first-order logic, developed by the German philosopher Gottlob Frege, assigns interpretation to some components of the language somewhat differently than does the Aristotelian tradition.[3]

Frege's first-order (predicate) logic provides a fairly general account of correct argumentation in mathematical reasoning. Because the reasoning in mathematics is possibly more accurate and better scrutinized than any other kind of reasoning, its formal model—first-order logic—became *the* logic for most scientists who pay attention to the formal details of argumentation. Because mathematics works with sharp definitions and its predicates are perfectly well characterized by their extensions, Frege constructed his logic (first-order logic) to be strictly extensional, even

[3]Aristotle thought that if "all men are mortal," then "some man is mortal" follows. For Frege, a universal sentence like "all men are mortal" is vacuously true if there are no men. The sentence would only be false if there was an immortal man.

though it was Frege (1892) himself who noted the limitations of extensionality. [4] In modern science and mathematics, first-order logic is the default; and it usually sits in the background. In other words, unless an analyst announces reliance on some other logic, we can safely assume that the logic to be used in analyzing the arguments is first-order logic. Moreover, our experience in the discipline leads us to think that first-order logic is the default in sociology as much as in any other science.

However it is also our experience that sociologists resist the suggestion that the patterns of reasoning in sociological theories are based on any particular logic (such as predicate logic). Our discussions with leading theorists suggest to us that they do not think that one or more logics belong in the toolkit of the discipline. Nonetheless, when questions arise about what conclusions follow from certain premises, these same analysts routinely appeal to the rules of "logic" (considerations of contradiction among premises, invalid inference, mistaken reasoning, etc.). Interestingly, appeals to logic enter the picture especially when someone contends the validity of a chain of argument, as in "Comments" and "Replies" sometimes published in the sociological journals.[5]

We conclude that theory building and reasoning in sociology do indeed depend on a logic but that this dependence is typically left implicit (until someone claims to find a gap or a contradiction in an argument). Because the explicit use of logic is episodic and informal, analysts have not generally made clear what kind of logics they deploy in analyzing the reasoning that underlies sociological theories. Perhaps the reason for this state of affairs is that analysts are tacitly assuming that there is only one kind of logic. Indeed, discussions in the social sciences about inference in arguments—what follows from what—generally make this tacit assumption. This is understandable because we generally learned logic in the context of studying mathematics, whose standard branches use only classical first-order logic.[6] This logic provides a formal characterization of the sound reasoning patterns typically found in mathematical analysis. Given that general education does not go beyond this logic, this is the end of the story for most social scientists.

But, in the discipline of logic, the story does not end here. Patterns of what seems to be sound argumentation can be found beyond the realm of mathematics, and some of these patterns do not fit into classical logics such as predicate logic. During most of the twentieth century, logicians were busy offering formal renderings for these other kinds of inference patterns. In the process, they created a very large number of logics.

A century of research in logic yields an important insight: what a set of premises implies can differ, depending upon what logic is chosen. As the leading logician Johan van Benthem (1996: 22) put it,

[4] Interestingly, there are still numerous applications for which only a fragment of this logic, such as the modern propositional logic, is needed.

[5] Not surprisingly, (predicate) logic plays a more central role in explicit treatments of theory construction, e.g., Coleman (1964), Stinchcombe (1968, 2005).

[6] Some parts of mathematics use other logics such as intuitionist logic and second-order logic. However, these are not the parts of mathematics that have been applied in sociology.

> Observed inferential patterns which seem "wrong" according to one notion of inference might just as well signal that the speaker is engaged in correct execution of patterns of another style of reasoning.

This book takes the implications of this view seriously. Indeed, van Benthem's claim could serve well as a motto for our effort to design a formal language for sociological theory.

The received wisdom of late-twentieth century logic indicates that *the choice of a logic* is a key step in building and integrating theories. Predicate logic initially seems to be an attractive choice. It is standard and familiar as one of the best-studied logics; and it also has many desirable properties.[7] This logic might provide the most appropriate tools for formal renderings of argumentation within established theories (but see below).

Unfortunately, using first-order logic requires more than sociology and other social sciences can usually deliver. Claims in sociological theories are generally *partial*. Instead of supplying the strict (universal) rules required for analysis in predicate logic, sociological arguments typically offer rules that admit exceptions. Such rules provide incomplete or partial accounts of the sociological regularities. Considerations of partial knowledge suggest that classical logics, such as predicate logic, might not suit the challenges of theory building in sociology.

In particular, predicate logic does not offer promise for the theory building and unification project that we undertake. We provide examples in the following chapters of the coexistence in organization theory of seemingly inconsistent theories. Making sense of the state of knowledge in such cases requires efforts to learn what claims can be unified and how the unification might best be attempted. Unfortunately, first-order logic offers only the following unsatisfying recipe for unification: collect all of the relevant premises and consider the consistency of the set. In the best case, the complete set of premises is consistent (in the sense that some state of the world could possibly satisfy all of the assumptions). In this happy case, the unification is complete. With less luck, we end up with an inconsistent set of premises, a failed unified theory.[8] Under these circumstances, it becomes obvious that some of the premises have to be dropped in the interest of consistency.

Unfortunately, predicate logic does not offer advice about what premises should be dropped. In other words, a set of partial theories (which we call *theory fragments*) stated as universal rules and formalized in predicate logic can rarely, if ever, be integrated. The key problems arise from the fact that this standard logic does not provide any way to isolate or control the different premises so as to avoid clashes among the universal rules from the different fragments.

A more useful strategy would be to choose a logic that provides some specific clues about how to deal with premises that lead to opposing predictions (inconsistencies). More generally, it would be better to choose (or define) a formal language that more accurately expresses how sociological theorists actually argue. We are

[7] These advantages include such properties as finite axiomatizability and compactness.

[8] Even if the unification does not yield inconsistency, it might be the case that the full set of premises imply unwanted theorems, theorems that run counter to the intuitions that motivate the theories, as we noted above.

convinced that theory building and theoretical argumentation in sociology do not adhere closely to the strictures of classical logics such as predicate logic (or the propositional fragment of predicate logic).

Sociologists typically do not often want to claim complete generality (universality) for propositions. They build theory fragments. As we use the term, a theory fragment consists of a set of interrelated assumptions, concepts, and propositions focused on one or several closely related explananda. Theory fragments typically arise from some core insight or intuition about these phenomena.

Upon scrutiny, sociology's theory fragments appear to be inconsistent, if the premises in each fragment are considered to hold universally. And they often rely (implicitly) on the specificity of arguments to control possible contradictions. We contend that casting arguments and patterns of reasoning in a dynamic logic provides a way to formalize sociological practice.

Moreover, premises in sociological arguments are not restricted to deal with facts. They often treat the "attitudes" of agents to some factual state of affairs, where attitude can mean perception, belief, valuation, and so forth. Such claims resist consistent translation into predicate logic (and other classical logics).

Perhaps the reason for sociologists' reluctance to admit that their theoretical argumentation relies on a logic is an implicit recognition that the language of classical logics does not fit the patterns of sociological argumentation. We agree about the mismatch. We think that the situation calls for examining the patterns of argumentation and the search for systematic methods for judging the soundness of arguments that do not fit classical patterns. This kind of effort might yield languages that allow sociologists and other social scientists to represent formally the actual structure of their arguments.

We try to show (without presenting all of the technical details) that dynamic logics hold promise for this task. (Chapters 6 and 7 introduce the particular dynamic logic that we use.) As we noted above, the generally used notion of logical inference holds that an inference is valid if the truth of the premises guarantees the truth of the conclusion. A "sound" logic satisfies this principle. First-order logic, along with most formal logics, builds on this principle.

An alternative way to provide a foundation to the idea of sound inference takes account of the *information content* of the premises and the conclusion. According to this view, if learning the premises (in a precise abstract sense) brings about the knowledge of the conclusion, then inferring the conclusion from the premises is logically valid. This dynamic notion of logical inference extends the usual static approach. When applied to the language of classical predicate logic, the static and dynamic approaches turn out to be equivalent. However, there is a broad class of situations in which the dynamic notion of inference applies but the classical static one does not.

In formalizing the theory-building process in dynamic logic, we refer to stages of a theory to characterize the information content of the theory (given the premises available at that stage). Claims of a stage of a theory tell how the world is expected to be according to the theory stage. Such claims need not be classical (universal) propositions; and inference based on these claims is best described by the dynamic approach to inference. These considerations motivated us to design a logic in line

with the dynamic approach.[9]

The variant of dynamic inference that we develop is a nonmonotonic logic. In a theory built on such a logic, claims can be stated as generic rules, that is, as rules-with-exceptions. (As we explain below, exceptions can be patterned, more than simply stochastic variation.) Specificity is used to control arguments, to decide which of a set of possibly applicable arguments should be used. Arguments are built of conditional sentences, sentences with an antecedent and a consequent, e.g., "If ϕ is the case, then ψ is the case." Now suppose we have two arguments that begin as follows: (1) "If an object is an organization, then . . . ;" and (2) "If an object is a young organization, then. . . . " Now consider the class of young organizations. It is clear that the antecedent in the second argument better approximates the class than does the antecedent in the first argument. In this case, we say that the second argument has greater specificity and is more relevant to the class of young organizations.

Establishing a specificity ordering is sometimes a straightforward matter unlikely to engender discussion, as in the foregoing example. In other cases, however, determining specificity orderings can be difficult and subject to varying interpretations. Indeed, it often takes considerable empirical knowledge of a subject to be able to settle matters of specificity.

The basic principle of nonmonotonic logic is that the most specific applicable argument prevails in the case of clashing implications. Our substantive applications seek to show that use of this nonmonotonic logic eases the task of integrating partly conflicting fragments while still allowing the theorist to derive the implications of an argument and to test the soundness of such inferences.

We supplement the nonmonotonic logic with a modal logic. Modern modal logics were developed to analyze arguments that use premises that express the attitudes (modalities) of agents. The issue of modality plays a central role in our theory of the interaction of producers and audiences in creating categories and forms.

Clearly, these logics are nonstandard tools. We justify their use by noting that they allow a formal treatment of sociological theories that expresses the actual patterns of argument used by leading theorists and researchers—not some disembodied notion about how theory should be constructed in an idealized world. Readers can judge for themselves whether this is the case for the class of organizational theories that we propose.

Stochastic Variations

When empirical researchers first confront the idea of rules-with-exceptions, they naturally think that this term refers to stochastic variations. Any stochastic process allows for exceptional events, events that occur with low probability given the structure of the process. There is no need for a new kind of logic to deal with such events—this is the business of probability theory.

Many processes of interest to social scientists are inherently stochastic. The formalizations presented in this book address stochastic variations explicitly. The key

[9] Translating this notion back to a standard truth-conditional approach is possible if truth is relativized to information states.

premises and conclusions hold for functions of probability distributions (usually expectations or hazards). For instance, a key outcome of interest is the hazard of organizational mortality, a function of the distribution of the lengths of lifetimes in a population of organizations.

What, then, do we mean by rules-with-exceptions in such cases? The exceptions refer to unusual probability distributions, those that might arise under specific circumstances. So, for instance, a theory might lead to the conclusion that the general case is such that the hazard of mortality of large organizations is lower than that of small ones. This is the relationship that we expect to find in normal circumstances, those in which more specific arguments do not come into play. Note that we do not expect that all large organizations will outlive all small ones. The usual stochastic variation in mortality processes makes it very unlikely that this will be the case. The exceptional cases that we have in mind are those in which specific considerations lead to a reversal of the general rule, cases in which small organizations persistently display a lower hazard.

Scope Restrictions

Science has long dealt with issues of partiality and possible contradictions among premises by restricting the scope of explanatory principles and causal claims.This is the only solution offered by classical logics, such as first-order logic, for dealing with clashing premises. If research identifies systematic exceptions to an argument, then theorists incorporate (strict) limits on the scope of applicability of premises. That is, they define precise bounds on processes such that the original argument holds within the bounds but not outside of them.

Reliance on scope restrictions to fix apparent inconsistencies dominates scientific practice in sociology and most social science. Arthur Stinchcombe (2005: 117) reflects the prevailing theoretical wisdom in sociology:

> Sociology is pervaded by boundary conditions on its causal processes . . . Contextual variations directly determine values of some causal or effect variables. They also determine the impact of variations in some causes on effects, and so produce different relations of causes to effects in one context than in another.

Under the strategy of relying on explicit restrictions on scope, the price to be paid for consistency is a limitation of explanatory power. Although we recognize the value of paying such a price, we insist that restricting the scope of premises is justified only if the restrictions can be well motivated substantively. Otherwise the cure is ad hoc, and it does not contribute to understanding.

Indeed, it takes a highly developed theoretical and empirical understanding to be able to state precise restrictions on the scope of arguments. One might well question whether sociologists and other social scientists can supply such precise information. Stinchcombe (2005: 121) clearly believes that they can; but he laments the current lack of attention paid to the problem: "we ordinarily want to spend as little time as possible on the theory of the context so as to spend as much as possible on the main subject of the research." He also recognizes the great difficulty of the task

of determining scope or boundary conditions:

> As a practical matter, many aspects of context vary a great deal from time to time and from place to place. Often these aspects are crucial for socially constructed boundaries, meanings and the various determinants of how sparse acts are. Our "second order units of analysis," on which we observe distances between contexts, frequently are times and places. The relations between times and places and social action are crucial to investigating distance between contexts. (Stinchcombe 2005: 1)

How should such boundary (or scope) conditions be expressed? So long as sociological arguments rely (often implicitly) on classical logic, there is no way to avoid stating these conditions as universals—at least implicitly. In other words, the implicit language of theorizing is such that scope conditions in the standard strategy do not admit any exceptions. In our experience, this is not what sociologists intend when they express limits on the scope of an argument. Knowledge about the limits on the scope of an argument is generally partial. This means that proposed scope conditions also generally come with not-yet-explained exceptions. Therefore, we conclude that the strategy of restricting scope does not really solve the problem of partiality if the pattern of reasoning is restricted to classical logic. We suggest that the use of a dynamic logic provides a way to handle limitations on scope formally in a way that better fits actual sociological practice.

To illustrate these points, consider an early contribution to organization theory. James Thompson, in his classic book, *Organizations in Action,* advanced many insightful propositions. One that influenced theory for many years states: "Under norms of rationality, organizations seek to buffer environmental influences by surrounding their technical cores with input and output components" (Thompson 1967: 20). Obviously, with this proposition (and many others in the book) Thompson explicitly stakes out a universal scope condition with the clause "under norms of rationality." This is admirable. Nonetheless, Thompson does not really do much with this scope restriction, either in terms of justification or implication in subsequent arguments. Instead, given its level of abstraction, the clause might be seen as a way to finesse around problematic empirical cases: should one encounter an organization that does not engage in buffering environmental influences with input and output components, then it can be said to operate under norms other than rationality. We think that similar usage surrounds much current practice where scope conditions are specified in sociological theory.

But now let us assume that the interpretation of the antecedent clause is not problematic and consider the rest of the proposition. It gives a "rule" for a structural process that organizations are expected to follow without exception. Research has shown clearly that very many organizations do follow this rule and exhibit the expected empirical implications. This was especially true in the 1970s and 1980s. However, with the advent of total quality management programs, just-in-time manufacturing practices, and other supply-chain management techniques driven by advances in information technology, it became clear in the 1990s that organizations did not need to buffer their technical cores with bureaucratic input and

output structures to be effective. Instead, a number of leading organizations showed that such structures were sometimes unnecessary and costly to maintain, that there were other, better ways to reduce uncertainty in the core by managing directly environmental agents. So, the rule does not appear to be universal—there seems to be good reason to think that some organizations might operate differently than it predicts, even when norms of rationality prevail. Unfortunately, the exact conditions under which certain organizations operate differently have yet to be clearly identified (at least to our knowledge).

What can a theorist do when facing this type of situation, assuming she wants to retain a degree of formalism and generality? Under the principles of first-order logic, the options are limited: either declare the proposition false or specify additional new scope conditions (stated as universals) knowing full well that these might be ad hoc and likely to need revision as soon as evidence appears. Both options are unattractive—the first because the original proposition does contain some insight and does explain many situations, the second because it misrepresents the state of theoretical knowledge.

Our impression is that sociologists and other social scientists typically respond to such situations by adopting the second option, while making it clear in their accompanying text that the formulation is tentative. In other words, they let it be known that the bounds of the argument are still unclear and that there are likely to be exceptions, which perhaps might be systematically delineated in the future. While perhaps sensible, this approach is easily subject to misinterpretation. In our opinion, the nonmonotonic logic that we put forward in this book does a similar job of specifying the state of theoretical knowledge but in a clearer, more systematic, and more parsimonious way.[10] As we explain below, this logic specifies propositions as rules with exceptions and uses differences in the specificities of arguments to determine which of several possible interpretations should apply. The textual markers that indicate we are employing nonmonotonic quantification in the equations consist of the words "normally" and "presumably." (The distinctions are also marked in the formal syntax, which uses the quantifiers $\mathfrak{N}$ and $\mathfrak{P}$ to express these kinds of quantification and also uses a third quantifier, $\mathfrak{A}$, to express nonmontonic quantification for auxiliary assumptions.)

1.3 PARTIAL MEMBERSHIPS: FUZZINESS

Classical logics such as the predicate (or first-order) logic—and, by extension, nearly all of mathematics—impose two important requirements. The first, discussed above, is that the theoretical statements must be universal, admitting no exceptions. The second requirement characterizes the concepts used. The concepts of a theory must satisfy what are typically called classical rules.

Frege developed logic in its modern form by (among many other things) making explicit the connections between logic and set theory, by linking concepts and sets.

[10] For a detailed and more complex illustration that parallels the Thompson example, see Chapter 10 on resource partitioning.

For Frege, the meaning of a concept is given by its extension, the set of objects for which the concept holds (in the sense that the statement that the object is an instance of the concept is true). For example, the meaning of the concept "red," according to this view, is given by the set of all those objects (in some specified universe of discourse) for which it is the case that they are red, that the statement "this object is red" is true. Knowing the membership of this set, the extension on the concept, is tantamount to understanding the concept. For this reason, the logics that build on the Fregian structure, such as modern predicate logic, are said to be *extensional.*

Frege sought to retain the classical principle that truth functions can take only two values: "true" and "false." Logicians refer to this principle as the *law of the excluded middle* (no truth value is allowed between false and true). Classical set theory fits this requirement, because a set is defined as a collection of objects for which a given property is true. Partial memberships in sets are excluded by definition. Frege (1893/1903, Vol. 2: 139) formulated his position on concepts as follows:

> The concept must have a sharp boundary. If we represent concepts in extension by areas on a plane, this is admittedly a picture that may be used only with caution, but here it can do us good service. To a concept without sharp boundary there would correspond an area that had not a sharp boundary-line all round, but in places just vaguely faded away into the background. This would not really be an area at all; and likewise a concept that is not sharply defined is wrongly termed a concept. Such quasi-conceptual constructions cannot be recognised as concepts by logic; it is impossible to lay down precise laws for them. The law of excluded middle is really just another form of the requirement that the concept should have a sharp boundary. Any object that you choose either falls under the concept or does not fall under it; *tertium non datur.* E.g., would the sentence "any square root of 9 is odd" have a comprehensible sense at all if *square root of 9* were not a concept with a sharp boundary? Has the question "Are we still Christians?" really got a sense, if it is indeterminate whom the predicate "Christian" can truly be ascribed to, and who must be refused it?

Frege's approach retains what E. E. Smith and D. L. Medin (1981) call the classical theory of concepts. Gregory Murphy (2002: 15) summarizes this perspective as follows:

> First, concepts are mentally represented as definitions. A definition provides characteristics that are (a) necessary and (b) jointly sufficient for membership in the category. Second, the classical theory argues that every object is either in or not-in the category, with no in-between cases . . . Third, the classical view does not make any distinction between category members. Anything that meets the definition is just as good a category member as anything else. (Aristotle emphasized this aspect of categories in particular.)

The classical theory has come under sustained attack in philosophy and cognitive science over the past half century. The leading figure on the philosophy side was Ludwig Wittgenstein. In his *Philosophical Investigations*, Wittgenstein (1953) repudiated his own influential work in logic (his *Tractatus*) and that of Frege and Bertrand Russell by abandoning the classical notion of the concept. His analysis of the social use of natural language led him to question whether our ordinary concepts satisfy the classical requirements and whether these requirements matter for human communication. Consider, for instance, his famous analysis of the concept "game."

> Consider for example the proceedings that we call "games." I mean board-games, card-games, ball-games, Olympic games, and so on. What is common to them all?—Don't say: "There *must* be something common, or they would not be called 'games'"—but *look and see* whether there is anything common to all.—For if you look at them you will not see something that is common to *all*, but similarities, relationships, and a whole series of them at that . . . And we can go through the many, many other groups of games in the same way; can see how similarities crop up and disappear . . .
>
> And the result of this examination is: we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail.
>
> I can think of no better expression to characterize these similarities than "family resemblances"; for the various resemblances between members of a family: build, features, colour of eyes, gait, temperament, etc. etc. overlap and criss-cross in the same way.—And I shall say: "games" form a family. . . .
>
> One might say that the concept "game" is a concept with blurred edges. —"But is a blurred concept a concept at all?"—Is an indistinct photograph a picture of a person at all? Is it even always an advantage to replace an indistinct picture by a sharp one? Isn't the indistinct one often exactly what we need? (Wittgenstein 1953)

Subsequent psychological research found broad and general support for the view that cultural concepts, coded in natural language, point to family resemblances rather than to something like classical concepts. The modern research tradition on the psychology of concepts began with Eleanor Rosch's (1973, 1975) examination of subjects' perceptions of the relationships of subconcepts to concepts. She asked subjects to tell how typical were certain fruits, e.g., apples, oranges, melons, olives, of the category "fruit." The subjects reported great differences in typicality among the specific kinds of fruit, and they agreed strongly on the typicality judgments. Apples, it was agreed, are very typical, and olives are very atypical. This pattern was replicated for "furniture" (tables and chairs are very typical items of furniture, and mirrors and carpets are atypical) and for "birds" (robins and sparrows are typical birds, and chickens and penguins are atypical). The robust finding that ordinary

concepts have an internal structure (of graded typicality) runs against the classical theory, which holds that all instances of a concept are alike.

Eleanor Rosch and C. B. Mervis (1975) claimed that these patterns reveal that concepts such as fruit and furniture constitute Wittgensteinian family resemblances. Several major lines of work in cognitive psychology and cognitive science have sought to explain the ubiquity of graded typicality for both natural concepts (such as those mentioned above) and artificial ones constructed for laboratory experiments. Although they disagree about the structure of categories, these strands of research agree broadly that concepts do not fit the classical story.[11]

One way to adapt to these empirical findings is by changing the foundational set theory. Standard set theory defines sets as the collections of objects in a universe of discourse that satisfy a given predicate, as noted above. The main alternative, fuzzy set theory, describes situations in which set membership can be partial, a matter of degree (Zadeh 1965). It defines a *grade-of-membership*, a function that maps from elements of the universe of discourse to the [0,1] interval. This function tells the degree to which the entity belongs to the set. Consider the grade-of-membership score for a single agent with regard to a particular concept. An agent's assigning an object a grade of membership of one indicates that the agent sees the object as clearly or unambiguously fitting the concept—there is no doubt about it. Conversely, a grade of membership of zero signifies that the agent has no doubt that the object does not fit the concept. For grades of membership between zero and one, the values tell the degree to which the agent categorizes the object in the concept.

It seems natural to link the structure of typicality found by Rosch and many others to vagueness of concept boundaries of the sort that worried Frege. Indeed, Rosch and other researchers made this connection in the 1970s. They suggested the general experimental finding that instances ("exemplars") of a category differ in typicality indicates that categories lack sharp boundaries. They proposed that fuzzy-set theory provided a language for analyzing this feature of concepts. In technical terms, this early research claimed a monotonic relationship between a grade of membership in typicality and a grade of membership as an instance of the concept (or category).

The initial enthusiasm for building a psychology of concepts on fuzzy-set theory was deflated by a series of papers that argued that concepts with fuzzy boundaries yield patterns that do not fit what we expect of a language. A basic property of a language is productivity, the idea that a speaker of a language can construct a new (not-yet-uttered) sentence that can be understood by the competent speakers of the language.

Productivity arises from *compositionality*, the idea, first formalized by Frege, that the meaning of a sentence ought to depend only on the meanings of the components and on the structure of the composition. The fuzzy representation of concepts can fail to conform to the principle of compositionality. In an influential pairs of papers, D. N. Osherson and E. E. Smith (1981, 1982) analyzed situations like the following. Suppose that we know what is a typical "pet" (a dog or cat) and what is a typical "fish" (a trout or salmon). What then do we expect of the concept "pet fish"? We

[11] Murphy (2002) gives an admirably lucid overview of these developments.

lack good grounds for providing an answer. Moreover, the typical "pet fish" (a guppy or goldfish) is a very atypical pet and also a very atypical fish. And here we have a potential problem with the use of fuzzy-set theory to represent typicality. In this set theory, the rule for constructing the intersection of two sets is that an entity's grade of membership in the intersection is the minimum of its grades of memberships in the two sets. The guppy or goldfish is problematic because it has a high grade of membership as a "pet fish" (the intersection) and low grade of membership in each of the component sets.[12]

The commentaries sparked by the Osherson-Smith critique made evident that a useful analytical distinction can be drawn between the two grades of membership (in typicality and in membership in the concept), between the internal structure of a concept (variations in typicality) and fuzzy boundaries (variations in membership)—see Laurence and Margolin (1999) and Hampton (1998). In other words, the existence of an internal structure to a concept does not imply the absence of a sharp boundary.

Cognitive psychologists responded to the critique that the use of fuzzy-set theory might sharply limit compositionality by drawing a sharp distinction between the notions of graded typicality (and the associated idea of an internal structure of concepts) and vagueness of category boundaries (Osherson and Smith 1997). Psychological research now downplays the idea of graded membership in concepts (vague concept boundaries)—but see James Hampton (2006) for an interesting exception. We think that this reaction went too far, at least as concerns socially constructed categories (as distinct from "natural kinds"). We discuss these issues in more detail in Chapter 2.

We follow the path blazed by Wittgenstein and followed by Rosch and her associates. We think that the social categories constructed by local audiences involve *both* typicality judgments and vagueness (or fuzziness) in category boundaries (Hampton 2006). Our reading of the evidence suggests that agents often detect shades of difference and decide that some producers fit comfortably in a category, some do not fit at all, and others fit to a greater or lesser degree. For example, consider the socially constructed organizational category "university" and some specific organizations that use—or at least sometimes lay claim to—the label. Everyone likely agrees that some obvious cases belong in the category: University of Bologna, Oxford University, Harvard University, and so forth. Then there are others which almost every audience member would say do not belong: the Hamburger University of McDonald's Corporation (an employee training center) is such a case, as is the Dreyer's Leadership University of the previous Dreyer's Grand Ice Cream Company.

Consider some more interesting cases. Rockefeller University is a great basic-

[12]Noun-noun combinations, such as "pet fish," allow wide scope for creativity in interpretation (Costello and Keane 2000). Some linguists argue that convergence in interpretations depends upon pragmatics, how the expressions are used. Moreover, it is worth noting that members of societies that lack the practice of keeping pet fish would not be expected to know that a guppy is a prototypical pet fish. Likewise, those unfamiliar with instance of various forms of financial institutions would likely have difficulty coming up with interpretations of such forms as "venture capital," "hedge fund," and "building society." See Werning, Machery, and Schurz (2005a, 2005b) for a variety of perspectives on the philosophical and empirical status of compositionality.

research organization but has few post-graduate students, no undergraduate students, and a tiny formal curriculum. The University of Phoenix calls itself the largest "university" in the United States in terms of student enrollment, has many online courses, has few regular or permanent faculty,and is part of a for-profit corporation. Bob Jones University has students and faculty, but also has a code of conduct subscribing to fundamental Baptist religion, including restraints on speech and research. The Maharishi University of Management has curricula in business and the humanities; and it emphasizes transcendental meditation by students and faculty as a central part of its academic mission. The National Defense University, which includes the National War College, trains top military and some State Department staff, gives masters degrees, and offers faculty appointments (that do not allow the possibility of tenure), and has a web page declaring that it subscribes fully to the AAUP guidelines on academic freedom. Britain's Open University enrolls 150,000 part-time long-distance students and does not require any qualifications for admission (other than a lower age limit). We suspect that much disagreement exists over the inclusion of these organizations in the university category. [13] Some might argue that one or another fits the category and others do not. More likely, they will claim that these organizations are universities "in some respects," are "technically" universities, or are "unusual" or "atypical" universities. The use of such hedge words is the classic linguistic sign of partiality—that is, fuzziness.

Even if social processes were to eliminate fuzziness for mature categories (contrary to what we see for "university"), it strikes us as implausible that fuzziness could be avoided at all stages of the emergence of a category as a cultural object. We try to account for the processes of emergence, beginning with efforts at clustering similar producers and products in social domains. We doubt that the conception of the emergence of graded typicality within a sharp concept (or category, as we call it) can serve as a credible story about emergence.

The theoretical research we report in this book treats the consequences of the categories created by audiences for various kinds of collective social action. The research in cognitive psychology that we have just sketched leads us to claim that these audiences create categories that lack classical properties. Fuzziness in category boundaries seems inescapable. Therefore, we seek to deal systematically with its implications head on.

This means that we confront two kinds of partiality: (1) incomplete knowledge about the causal processes at work; and (2) partial applicability of concepts and categories. As we see it, the first kind of partiality characterizes the situation facing the theorist, and the second kind characterizes the situation facing the agents whose actions the theory treats (and possibly also the theorist). We think the two kinds of partiality have different impacts on arguments, and we want to distinguish them clearly. Therefore, we mark the two kinds of partiality in our formal language. We represent the first kind of partiality as *genericity* (and we develop quantifiers that can deal with genericity). We represent the second kind of partiality by *graded membership* in categories. This means that the fuzziness is restricted to the lan-

[13] Interestingly, the Open University's Web site has a section entitled "Is the Open University a 'real university'?"

guage of the agents; our theoretical language is classical in this respect.[14]

1.4 ORGANIZATIONAL ECOLOGY

We next introduce the main subject of our theory building: the program of organizational ecology. This body of theory and research examines interactions within and between populations of organizations. The approach differs from other sociological research on organizations by focusing on the population level, relying on selection mechanisms, and studying the life histories of all organizations in populations over their full histories. Organizational ecology initially borrowed ideas from neoclassical population bioecology, which analyzes numerical aspects of population interactions from an evolutionary perspective. In paraphrasing the bioecologist G. Evelyn Hutchinson (1959), Michael Hannan and John Freeman (1977) posed the orienting question: Why are there so many kinds of organizations? They suggested that social science at the time lacked good answers to this question and that seeking answers would clarify the dynamics of the organizational world. Three decades of subsequent research shows clearly that organizational diversity exerts important consequences for individuals and social structures (Carroll and Hannan 2000).

Organizational ecology sought initially to resolve a theoretical tension about what sociologist Amos Hawley (1968: 334) termed the principle of isomorphism: "Units subjected to the same environmental conditions or to environmental conditions as mediated through a given key unit, acquire a similar form of organization." Hannan and Freeman (1977) argued that Hawley's principle, like other adaptationist premises, does not apply straightforwardly in uncertain, heterogeneous environments. Extending the received wisdom about adaptation requires specification of the underlying dynamic processes and attention to selection. In the organizational world, selection occurs through the emergence and demise of organizational forms, which depend on the fates of individual organizations. Accordingly, theory and research in this tradition focuses on the demographic vital rates of organizations and organizational populations: rates of founding, growth, and mortality.

The ecological research program progressed by fostering empirical research in a variety of distinct theory fragments. Among other topics, the major theory fragments of organizational ecology address questions about:

Organizational forms and populations. This fragment addresses questions about how to define forms and populations and how to classify them meaningfully into higher-order forms. Early ideas focused heavily on patterns of exchange among organizations and other environmental actors; however, recent theory and research centers on ideas about social identities and social codes. (See Hannan and Freeman 1986; Zuckerman 1999; McKendrick and Carroll 2001; Rao, Monin, and Durand 2003, 2005; Ruef 2000, 2004a; and Baron 2004.)

[14]This division of assignment contrasts markedly with that of some other applications of logic in sociological theory where the analytical problems faced by the theorist are characterized by fuzzy concepts (see Montgomery 2000; Ragin 2000).

Structural inertia and change. The inertia fragment develops arguments about the rigidity of organizational structures and argues that strong inertia makes selection an important motor of change in the world of organizations. It addresses the main possible mechanisms behind such phenomena, including the predilection in modern society to value accountability and reliability, as well as inertia's evolutionary implications. (See Hannan and Freeman 1984; Haveman 1992; Amburgey, Kelly, and Barnett 1993; Barnett and Carroll 1995; Baron, Hannan, and Burton 2001; Hannan, Pólos, and Carroll 2003a, 2003b, 2004; and Phillips and Owens 2004.)

Age dependence. Theory and research on age dependence asks how and why the age of organizations matters for their structures and life chances. The proposed answers to the problem, which involve issues such as knowledge, capabilities, bureaucratization, and obsolescence, transcend the seemingly narrow question. Yet theoretical progress in the fragment has been clouded by conflicting empirical evidence. (See Carroll 1983; Freeman, Carroll, and Hannan 1983; Levinthal 1991; Barron, West, and Hannan 1994; and Sørensen and Stuart 2000.)

Dynamics of social movements. Social movement research in organizational ecology emphasizes the organizational basis of collective action, especially that related to the competition and mutualism of movement organizations. It also ties movements to the rise of new organization forms. Social movement theorists naturally focus on the possibility that institutions can sometimes be changed, and they stress the importance of attending to movement audiences and their dynamics. (Relevant publications include Hannan and Freeman 1987; Minkoff 1999; Ingram and Simons 2000; Olzak and Uhrig 2001; Swaminathan and Wade 2001; Sandell 2001; Koopmans and Olzak 2004; and Greve, Posner, and Rao 2006.)

Density dependence. This theory fragment comprises what is perhaps ecology's most sustained research program on population dynamics, the model of density dependence in legitimation and competition. The core theory posits relationships between density, the number of organizations in a population, and legitimation of the form of organization and competition among the population's members. Its main empirical implications are nonmonotonic relationships between density, on the one hand, and population vital rates on the other hand. Extensions to the theory attempt (1) to extend the model to explain late-stage declines in population density, an observed empirical regularity, and (2) to treat legitimation as "sticky" or not easily reversible. (See Hannan and Freeman 1989; Carroll and Hannan 1989; Hannan and Carroll 1992; Barron 1999; and Ruef 2004b.)

Niche structure. An organization's niche summarizes its adaptive capacity over the various possible states of its environment. Theories in this fragment build on the concept of niche width, the span of environmental states in which an organization can thrive. These theories claim that a broad niche comes at the

expense of viability in a stable, competitive environment, but that environmental uncertainty and variability affect the tradeoff between niche width and viability. (See Hannan and Freeman 1977; McPherson 1983; Freeman and Hannan 1983; Baum and Singh 1994; Podolny, Stuart, and Hannan 1996, Dobrev, Kim, and Hannan 2001; Sørensen 2000; Dobrev, Kim, and Carroll 2003; and Barnett and Woywode 2004.)

Resource partitioning. This fragment can be seen as a variant of general niche theory, one based on different assumptions and scope conditions. This fragment explains the endogenous partitioning of markets (environments) as an outcome of competition between populations of generalists and specialists. (See Carroll 1985; Swaminathan 1995; Carroll and Swaminathan 2000; Péli and Nooteboom 1999; Park and Podolny 2000; and Boone, Carroll, and van Witteloostuijn 2002.)

Diversity of organizations. Research in this fragment deals with the social and economic consequences of the level of diversity among the types of organizations in a community or sector. An initial stream deals with the interplay of careers of individuals and the organizational ecologies within which careers play out. (See Hannan 1988; Carroll, Haveman, and Swaminathan 1990; Greve 1994; Haveman and Cohen 1994; Fujiwara-Greve and Greve 2000; Phillips 2001; and Sørensen and Sorenson 2007. Recent research addresses religiosity (Koçak and Carroll 2006) and election turnout (Carroll, Xu, and Koçak 2005).)

As these brief summaries illustrate, organizational ecology contains diverse theory fragments and associated lines of empirical inquiry. Moreover, these fragments can sensibly be regarded as parts of a larger research program—they build on a common conception of the organizational world as shaped by processes of selection operating on organizational forms and also share methodological presumptions and practices (Carroll and Hannan 2000). Nonetheless, the relationships of the fragments to each other remain ambiguous. Key points of apparent conceptual intersection sit unexplored and require clarification. The preponderance of effort over the last three decades has focused on empirical testing, with less attention paid to issues of theoretical integration. This empirical emphasis, though highly successful for producing new knowledge, limits further progress. Because empirical research has not been balanced by theoretical efforts to examine the relations among the fragments, we lack a clear vision of what empirical projects would move the larger program forward substantially.

An effort to integrate the fragments provides an opportunity to rethink foundational issues in the light of the successful record of research. Initial theoretical formulations did not have a strong empirical base. Some aspects of these formulations bore fruit; others did not. Naturally, many processes turned out to be more complicated than anticipated in the early phases of theoretical work. Accordingly, we do not think it worthwhile to try to fit all of the fragments into a single whole or to insist on the priority of the initial theoretical claims. Instead, we try to retain what we see is a core set of apparently related insights that proved to be useful in

the various facets of the research program.

This theory-building strategy leads us to place social codes at the forefront of ecological analysis. As we use the concept here, a social code denotes and connotes both cognitive recognition and imperative standing. A social code can be understood (1) as a set of interpretative signals, as in the "genetic code," and (2) as a set of rules of conduct, as in the "penal code." Some of the most interesting and important processes are those that convert interpretative schemata into imperative codes.

Our theoretical strategy of rebuilding foundations so as to enable integration stands in sharp contrast to one that animated an earlier vigorous phase of formalization of fragments of organizational ecology. In this earlier phase, initial formulations of various fragments were subjected to rational reconstruction and logical analysis designed to test the soundness of the arguments. Fragments analyzed in this manner include (1) structural inertia and change (Péli et al. 1994; Péli, Pólos, and Hannan 2000); (2) niche width (Péli 1997); (3) life-history strategies (Péli and Masuch 1997); and (4) age dependence (Hannan 1998; Pólos and Hannan 2002). These efforts took seriously the "frozen" published texts. They translated the natural-language renderings of the arguments into a formal language (and sought to tune translations to the intuitions that animated the original arguments) and checked the proofs of the claims in that language. This work has been valuable in establishing the soundness of the main arguments and in filling gaps in arguments. However, the approach takes a largely passive role with respect to moving the theories forward. In particular, because these efforts considered each fragment in isolation from the others, they could not clarify the relationships among the fragments.

We take a much more active stance in this book. We seek to rebuild, formalize, and integrate the fragments while still trying to preserve the main substantive insights. As the Table of Contents indicates, our efforts mainly encompass the fragments about age dependence, forms and populations, niche structure and resource partitioning, and structural inertia and change.

1.5 UNIFICATION PROJECTS

Fragmentation in a discipline or a theory program (as just sketched for organizational ecology) seems to be endemic in sociology. If this assessment is accurate, then developing strategies for unification ought to be high on the agenda of sociological methodology (in the broad sense).

Fragmentation clearly is an outgrowth of the common mode of tying empirical investigation to the development of the type of middle-range theory advocated by Robert K. Merton. In responding to the grand sociological theories of Talcott Parsons and others, Merton (1968) proposed that greater scientific progress could be achieved if sociological theory focused on problems in the middle range. He described middle-range theories as those

> . . . that lie between the minor but necessary working hypotheses that

> evolve in abundance during the day-to-day research and the all-inclusive systematic efforts to develop a unified theory that will explain all the observed uniformities of social behavior, social organization and social change. (Merton 1968: 39)

Merton argued that despite their specific foci, middle-range theories also possess a "general character" that allows them to be applied in new contexts and elaborated conceptually. Yet

> . . . it is equally clear that such middle-range theories have not been logically *derived* from a single all-embracing theory of social systems, though once developed they may be consistent with one. Furthermore, each theory is more than a mere empirical generalization—an isolated proposition summarizing observed uniformities of two or more variables. A theory comprises a set of assumptions from which empirical generalizations have themselves been derived. (Merton 1968: 41)

Although social scientists do not commonly use the term, what we call a theory fragment bears a close resemblance to the Mertonian notion of a middle-range theory.

In our view, the strategy of working at the middle range mitigates some serious obstacles to conceptual innovation. It eases the tasks of formulating and communicating new assumptions, propositions, and the like. In particular, new fragments of middle-range theory do not face the possibly formidable challenge of integrating tightly with existing theory. Accordingly, the strategy of constructing middle-range theory allows new insights to develop rapidly and diffuse through the scientific community—spurring a healthy development of scientific knowledge through the flowering of ideas.

Yet, a bouquet of theoretical blossoms yields a certain theoretical fragmentation that limits the broader development and use of a body of scientific knowledge. These issues come into sharp focus when we consider prediction and empirical testing, cornerstones of any scientific knowledge base. Fragmentation raises a serious problem for prediction: different theory fragments might yield different predictions for a not-yet-seen case. What, then, is the predictive power of the larger body of knowledge in which the fragments sit? A similar problem arises in empirical testing. With fragmentation, the relation of a given empirical finding to a set of theory fragments is generally indeterminate (so long as the interrelations of the fragments are problematic); therefore, empirical implication has smaller scope under fragmentation.

A profusion of theories of the middle range can be likened to the consequences of fertilizing a garden. Initially, new and old vegetation of a wide variety sprouts and blossoms; the garden looks very lush because of all the new growth. After a period of time, however, it needs pruning to remain healthy. Some plants grow too big and block the sun. Dead and diseased plants need to be removed, and vibrant ones need to have some space cleared for further growth.

In similar fashion, a new burst of middle-range theory construction in a substantive area eventually needs to have some order restored through integration. An

integration project might cause some fragments to be discarded (or subsumed by others), while other fragments are straightened out and given more attention. To be more specific, the challenge of dealing with fragmentation involves learning (1) if and how the various fragments fit together, and (2) what changes must be made to some fragments to make them fit (and whether these changes retain the basic insights). In fact, some changes usually must be made, because fragments are typically tuned to particular applications for which specialized assumptions might be warranted. However, theoretical integration should embrace the goal of preserving the core insights whenever possible (especially when the fragments have survived empirical testing) while making the fragments fit together so that their interrelations can be stated precisely.

Theory fragments come with different sets of concepts and assumptions. Attempting to unify disjoint conceptual vocabularies is uninteresting, because nothing can be done. However, if the sets of concepts overlap, then it becomes interesting to investigate the compatibility of the assumptions of the different fragments.

We do not propose that all social science theory would currently benefit from attempts at unification. If serious doubts arise about the arguments in some available fragments, unification can produce disastrous results. Consider, for instance, a unification of a pair of hypothetical theory fragments A and B, where B happens to be very misguided (in hindsight), possibly based on flimsy or incorrect empirical evidence. The problem, of course, is that the unified theory is also necessarily flawed and, in the worst-case scenario, so too are the remains of fragment A, which have now become intertwined with B as it developed.

Theories often emerge as the results of efforts more or less coordinated by research programs conducted by many individual scientists. This kind of origin does not guarantee that different fragments of a theory use concepts in a coordinated manner. What does it take to develop a coordinated conceptual inventory? First comes the task of eliminating ambiguities and inconsistencies of different usages of concepts. As long as the concepts sit close to the surface of the theory, i.e., the key concepts are explained informally but not defined in terms of other primitive terms, then concepts can be changed easily, and the changes remain (relatively) local. For this reason, it makes sense to use locally defined (close-to-the-surface) conceptual inventories in the early phases of the development of a theory. The absence of deep, intricate connections between different concepts carries a less desirable side effect, though. It limits the explanatory and predictive power of the theory. If, on the other hand, the undefined, atomic concepts of a theory lie further from the surface (i.e., the important theoretical concepts are defined via sequences of definitions), then the modular nature of the conceptual inventory leads to more insightful explanations and more informative predictions. All these benefits become increasingly important as the theory matures.

Building deeper definitional structures requires an exploration process. Analysis of the "intended content" of the surface concepts leads to the deeper definitions and meaning postulates. Generally such an exercise can be carried out in different ways. One might go deeper, while another might try to stay closer to the surface. Some disassemble a complex notion in one way, and someone else does it differently. (A square can be seen as a rectangular rhombus or as an equilateral rectangle.)

Even though there is no single best way to execute these theoretical excavations, a number of indicators signal whether a process is running on a good track. If the number of atomic (undefined) concepts stops growing, (i.e., new definitions reuse the already excavated atomic concepts), then the network of definitional relationships gets denser. This is a good sign! If, as a result of such explorative work, previously unrelated theory fragments now jointly imply new consequences, then the conceptual excavation starts to reveal treasures. On the other hand, if deeper definitional structures do not bring unforeseen new connections, and if the new theorems derivable as a consequence of adding the new definitions show the signs of triviality, then one might seriously consider stopping the dig.

Attempts at unification presuppose a certain level of confidence in the component fragments. Such confidence will always reflect judgment; but it seems safe to say that attempting to integrate new fragments, or those without much empirical basis, typically involves greater risk. Such fragments might be better left isolated or unintegrated, at least until they mature as fragments with strong empirical foundations.[15]

As long as analysts restrict themselves to the rules of classical (first-order) logic as the way to test the soundness of arguments, the dangers of integration are still greater if two or more fragments offer incompatible explanations of a phenomenon. Suppose, for instance, that fragment A explains a set of facts from one set of premises and fragment B explains the same facts from incompatible[16] premises. The incompatibility means that the unification of these fragments will be logically inconsistent. Inconsistency is a theoretical disaster. The standard approach to testing the soundness of an argument is to analyze the set of premises together with the *negation* of the purported conclusion. If this set can be shown to be unsatisfiable (in the sense that there is no set of circumstances that make all of the premises and the negation of the conclusion true), then the theorem is proven, the conclusion follows as a logical implication of the set of premises (see Appendix C). This leads to a surprising conclusion: an inconsistent theory would explain all the facts that any fragment explained, and it would also explain a lot more, even the negations of all of these facts. Indeed, every possible theorem can be shown to follow as an implication of an inconsistent premise set.

The parallel to inconsistency in a classical logic shows up in the dynamic logic that we propose as the presence of arguments of equal or incomparable specificity lead to opposing conclusions. Any theory that contains such arguments is not much more useful in our (nonclassical) framework than it is in the classical framework. Instead of implying everything, opposing arguments with equal or incomparable specificity block any conclusion (see Chapter 6).

Given what seems to be the common interpretation of the strategy of building middle-range theory, it might come as a surprise to learn that the type of unification that we advocate does not contradict Merton, his views on middle-range theory

[15] In this book, we sometimes end chapters (and treatments of various fragments) by making speculative theoretical claims that seem to be natural extensions of developed arguments but for which we have no empirical evidence. Although these extensions might be readily made formal and integrated into the developed body of theory, we usually refrain from doing so for this reason.

[16] Incompatibility means that it is impossible for both sets of premises to be true simultaneously.

notwithstanding. In describing his full views about how theoretical activity should proceed, Merton wrote:

> Sociological theory, if it is to advance significantly, must proceed on these interconnected planes: (1) by developing special theories from which to derive hypotheses that can be empirically investigated [that is, with middle range theories] and (2) by evolving, not suddenly revealing, a progressively more general consolidated group of special theories. To concentrate entirely on special theories is to risk emerging with specific hypotheses that account for limited aspects of social behavior, organization and change that remain mutually inconsistent. (Merton 1968: 51)

So, although it seems to have been forgotten or disregarded by many,[17] there is no question that Merton saw middle-range theory and research as only an initial step in a larger process that eventually produces a more unified theory. This book deals with the subsequent step: integration.

Given the potential benefits of theoretical integration, why does scientific activity of this kind play such a limited role in sociology and other social sciences? One answer is that many researchers value integration but regard this activity as little more than collecting disparate fragments together into an available portfolio of theories (rather than pruning, trimming, and discarding some of them).[18]

Another answer points to social processes: theory groups, collections of analysts who support a theory, form around fragments. These groups commonly see themselves as advocates of the theory, and they commonly resist attempts to weaken its identity and prominence, as integration might possibly do. Interestingly, the view of integration as appending fragments with minimal alteration dovetails nicely with the interests of theory groups because, in this approach, the fragments retain their original coherence.

A third answer to the question about the rarity of integration projects recognizes that, even if a social scientist believes in deep integration and acts dispassionately towards the theory fragments of interest, it would often be hard to make progress. This is because good tools for theoretical integration are largely absent from the social-science toolkit for theory building, leaving the impression that integration relies more on genius and deep insights rather than on systematic procedures. We hope that our efforts to use a formal logic as a method for theoretical integration in rebuilding and extending a line of sociological theory will demonstrate that systematic procedures can be developed (and, to some extent, already exist implicitly in the best practice of contemporary social scientists).

Finally, we also suspect that the uncertainty of knowing where and how to start an integration project plays a role. If various fragments appear to be connected to others but in unclear ways, it is hard to pick the place where foundational concepts and arguments should be developed a priori. We have no good solution to this

[17] For example, Stephen Cole (2001), a student of Merton, laments the use of abstract general theory in sociology, suggesting (or so it seems) that only middle-range theory building makes sense.

[18] Some claimed integrations amount to little more than including indicators associated with causal processes featured in different theories as covariates in some kind of lengthy regression equation.

dilemma; we can only suggest that trial-and-error makes it possible to discover eventually an integration that proves useful.

In fact, our own efforts experienced considerable backtracking and rethinking as we attempted to integrate fragments. As our record of previous journal publication illustrates, we began by addressing what we saw as foundational issues concerning organizational forms and structural inertia. We rebuilt theories of each, using what we saw as the core concepts: audiences, social codes, and identity. We then proceeded to redraw theories of ecological niches and resource partitioning using this framework. As we then attempted to prepare a fuller treatment in this book, we realized that it was important to develop a theory of form emergence. After great deliberation, our work on form emergence eventually caused us to make the agents more explicit, attend to their perceptions, and emphasize the degree of agreement among the agents. This approach then led us to introduce vagueness in perception and grades of membership in clusters, categories, and forms—developments that begged us to deploy fuzzy-set theory and perception and default modalities. However, after doing all this, we realized that much of our previously reconstructed fragments were obsolete: they did not incorporate the very basic notion of vagueness in perception but instead something that denied this key insight. So we then reconstructed the reconstructions, building these ideas into the various fragments, sometimes with great pain but also at other times with interesting new insights. We report the results of this iterative effort in this book.